

IBM Tivoli Storage Productivity Center  
Version 5.2.4

*User's Guide*



**Note:**

Before using this information and the product it supports, read the information in "Notices" on page 1149.

This edition applies to version 5, release 2, modification 4 of IBM Tivoli Storage Productivity Center (product numbers 5725-F92, 5725-F93, and 5725-G33) and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC27-4060-04.

© **Copyright IBM Corporation 2001, 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

## About this guide. . . . . ix

Who should read this guide . . . . .	ix
Publications . . . . .	ix
Accessing publications online . . . . .	ix
IBM Redbooks . . . . .	xi
Translation. . . . .	xi
Printable documentation . . . . .	xii
Providing feedback about publications . . . . .	xiii
Tivoli Storage Productivity Center Service Management Connect community . . . . .	xiii
Contacting IBM Software Support . . . . .	xiii
Reporting a problem . . . . .	xiv
Conventions used in this guide . . . . .	xiv

## | New for Tivoli Storage Productivity Center Version 5.2.4 . . . . . xvii

### Chapter 1. Getting started . . . . . 1

Managing resources . . . . .	6
Configuring and administering. . . . .	10
Navigation . . . . .	11

### Chapter 2. Product licenses. . . . . 17

Functions available in the web-based GUI based on role and license . . . . .	19
Functions available in the stand-alone GUI based on role and license . . . . .	24

### Chapter 3. Product overview . . . . . 31

Architecture . . . . .	31
Interfaces for Tivoli Storage Productivity Center . . . . .	33
Available functions in the interfaces . . . . .	34
Functions that were moved to the web-based GUI in V5.2 and later . . . . .	39
Starting Tivoli Storage Productivity Center . . . . .	42
Starting the Tivoli Storage Productivity Center stand-alone GUI. . . . .	42
Starting the Tivoli Storage Productivity Center web-based GUI . . . . .	44
Starting Tivoli Storage Productivity Center for Replication . . . . .	45
Tivoli Storage Productivity Center for Replication . . . . .	46
Key components of Tivoli Storage Productivity Center for Replication . . . . .	48
Interfaces for Tivoli Storage Productivity Center for Replication . . . . .	49
Tivoli Storage Productivity Center for Replication integration with Tivoli System Automation . . . . .	52
Application Manager . . . . .	52
Management server. . . . .	52
Storage systems . . . . .	52
Host systems . . . . .	52
Ports . . . . .	53
Sessions . . . . .	55

Session types . . . . .	57
Metro Mirror heartbeat . . . . .	64
Site awareness . . . . .	64
Users and groups . . . . .	64
User roles . . . . .	64
Key concepts . . . . .	65
Data collection . . . . .	65
Performance monitoring and troubleshooting . . . . .	66
Applications and departments . . . . .	68
Storage optimization . . . . .	69
Cloud configuration . . . . .	70
Provisioning storage . . . . .	75
Tivoli Storage Productivity Center VASA provider . . . . .	78
vSphere Web Client extension for Tivoli Storage Productivity Center. . . . .	79
Alerts . . . . .	80
Reporting . . . . .	82
Units of measurement for storage data . . . . .	82
Role-based authorization . . . . .	83
Fabrics and zones . . . . .	85
Agents . . . . .	88
Single sign-on . . . . .	90
Launch and launch in context . . . . .	92

### Chapter 4. Scenarios. . . . . 93

Optimizing the performance of storage virtualizers . . . . .	93
Redistributing volume workloads to balance pools . . . . .	94
Re-tiering volumes to balance pools . . . . .	95
Tiering volumes by I/O density and file usage . . . . .	97
Tiering volumes by I/O rate. . . . .	99
Collocating volumes . . . . .	101
Monitoring capacity usage at different levels of a business hierarchy. . . . .	102
Comparing storage usage in each department . . . . .	104
Using applications and subcomponents to monitor capacity and space usage . . . . .	106
Viewing storage capacity and usage trends . . . . .	108
Tracking performance issues on a storage system . . . . .	110
Scheduling data collection . . . . .	110
Defining a performance threshold . . . . .	110
Viewing threshold violations . . . . .	111
Limiting a set of storage to provision in the vSphere Web Client extension . . . . .	112
Viewing NPIV connections between server ports and switch ports in a fabric. . . . .	113
Replication scenarios . . . . .	114
Performing planned maintenance scenarios . . . . .	114
Recovering from a disaster scenarios . . . . .	117
Practicing disaster recovery scenarios . . . . .	120

### Chapter 5. Managing resources . . . . . 127

Adding resources . . . . .	127
----------------------------	-----

Resources that you can monitor in the web-based GUI . . . . .	128	Configuring Tivoli Storage Productivity Center for provisioning . . . . .	245
Adding storage systems . . . . .	132	Setting the zoning policy . . . . .	261
Adding fabrics and switches . . . . .	133	Provisioning storage with the Tivoli Storage Productivity Center GUI. . . . .	262
Adding hypervisors . . . . .	135	Provisioning storage with the vSphere Web Client extension . . . . .	266
Adding servers . . . . .	138	Optimizing storage tiering . . . . .	269
Removing resources . . . . .	144	Setting the tier level of storage pools . . . . .	270
Removing storage systems . . . . .	145	Analyzing tiering by servers . . . . .	271
Collecting data . . . . .	145	Analyzing tiering by hypervisors . . . . .	271
Collecting asset and status data . . . . .	147	Analyzing tiering by storage virtualizers . . . . .	272
Collecting performance data . . . . .	151	Analyzing tiering by storage pools . . . . .	272
Collecting information about shares on storage systems . . . . .	159	Analyzing tiering by volumes . . . . .	272
Alerting . . . . .	159	Balancing pools and re-tiering volumes. . . . .	273
Alert severities . . . . .	161	Tiering volumes by I/O density, I/O rate, and file usage. . . . .	273
Defining alerts . . . . .	162	Modifying the criteria for analyzing tiering . . . . .	274
Viewing alert definitions . . . . .	170	Storage tiering policies . . . . .	274
Viewing and modifying alert notification settings . . . . .	174	Optimizing storage pools . . . . .	276
Disabling or enabling alert definitions . . . . .	178	Balancing pools . . . . .	276
Modifying alert definitions . . . . .	178	Modifying the criteria for balancing pools . . . . .	278
Renaming alert definitions . . . . .	180	Criteria for identifying the pools that require balancing. . . . .	278
Viewing all alerts . . . . .	181	Transforming and migrating volumes . . . . .	280
Viewing alerts for a type of resource . . . . .	183	Modifying the period for analyzing performance data . . . . .	281
Acknowledging alerts . . . . .	184	Modifying the properties of resources . . . . .	281
Removing alerts . . . . .	184	Opening the management GUI for a resource. . . . .	284
Viewing information about resources . . . . .	185	Opening the DS Storage Manager GUI for a DS8000 . . . . .	285
How information is organized. . . . .	185	Exporting information about resources, jobs, and alerts . . . . .	285
Dashboard view . . . . .	188	Customizing lists of resources, tasks, and alerts . . . . .	287
Viewing information about top-level resources	193	Filtering lists . . . . .	287
Viewing information about internal and related resources . . . . .	194	Sorting lists . . . . .	288
Display differences in the interfaces . . . . .	195	Showing, hiding, and reordering columns in lists . . . . .	289
Monitoring the status of resources . . . . .	196	Managing tasks. . . . .	289
How status is determined . . . . .	197	Viewing tasks . . . . .	289
Viewing the overall status of resources . . . . .	200	Managing tasks for provisioning . . . . .	293
Viewing the status of top-level resources . . . . .	201	Managing tasks for tiering storage, balancing pools, and transforming storage . . . . .	295
Viewing the status of internal resources . . . . .	201	Renaming tasks . . . . .	301
Acknowledging the status of resources . . . . .	202	Deleting tasks . . . . .	301
Monitoring the performance of resources . . . . .	205	<b>Chapter 6. Managing replication . . . . . 303</b>	
Viewing performance information . . . . .	206	Administering . . . . .	303
Viewing threshold violations and alerts. . . . .	218	Starting and stopping Tivoli Storage Productivity Center for Replication . . . . .	303
Monitoring and administering applications and departments . . . . .	221	Verifying that the Tivoli Storage Productivity Center for Replication server is running . . . . .	304
Applications. . . . .	225	Identifying the version of Tivoli Storage Productivity Center for Replication . . . . .	305
Departments . . . . .	233	Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data. . . . .	305
Exporting information about applications and departments. . . . .	236	Verifying the Java heap size . . . . .	307
Viewing the data paths of resources . . . . .	238	Managing management servers . . . . .	308
Resources that you can view in a data path . . . . .	239	Management servers . . . . .	308
Navigating in a data path . . . . .	242	SNMP alerts. . . . .	310
Viewing information about top-level resources	242		
Viewing information about internal resources	242		
Zooming in . . . . .	243		
Zooming out . . . . .	243		
Changing the layout . . . . .	243		
Saving the display settings . . . . .	244		
Printing a view. . . . .	244		
Exporting a data path to a file. . . . .	244		
Provisioning storage . . . . .	244		

Setting up a standby management server . . . . .	313	Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session . . . . .	436
Applying maintenance to an active management server . . . . .	314	Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session . . . . .	436
Reconnecting the active and standby management servers . . . . .	315	Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session . . . . .	437
Performing a takeover on the standby management server . . . . .	315	Monitoring health and status . . . . .	437
Configuring SNMP . . . . .	316	Viewing the health summary . . . . .	437
Adding SNMP managers . . . . .	316	Viewing SNMP alerts. . . . .	438
Changing the standby management server port number . . . . .	316	Viewing sessions . . . . .	438
Changing the client port number. . . . .	317	Viewing session properties . . . . .	447
Managing storage systems . . . . .	317	Viewing copy sets in a session. . . . .	476
Storage systems . . . . .	317	Viewing storage system details . . . . .	476
Storage connections . . . . .	319	Viewing storage connection details . . . . .	476
Protected volumes. . . . .	323	Viewing volume details . . . . .	476
Site awareness . . . . .	323	Viewing logical paths. . . . .	477
Adding a storage connection . . . . .	324	Viewing host connection details . . . . .	477
Removing a storage system. . . . .	327	Viewing console messages . . . . .	477
Modifying the location of storage systems. . . . .	328	Security . . . . .	477
Modifying storage connection properties . . . . .	328	Users and groups . . . . .	478
Refreshing the storage system configuration . . . . .	329	User roles . . . . .	479
Setting volume protection . . . . .	329	Granting access privileges for a user. . . . .	480
Restoring data from a journal volume . . . . .	330	Viewing access privileges for a user . . . . .	481
Configuring the use of the Easy Tier heat map transfer function . . . . .	330	Modifying access privileges for a user . . . . .	481
Managing host systems . . . . .	333	Removing access privileges for a user . . . . .	482
Adding a host system connection. . . . .	333	<b>Chapter 7. Reporting . . . . . 483</b>	
Modifying a host system connection. . . . .	334	Getting started . . . . .	483
Removing a host system connection . . . . .	335	Predefined reports. . . . .	484
Removing a session from an AIX host system connection . . . . .	336	Predefined reports listed by package . . . . .	485
Managing logical paths . . . . .	336	Custom reports. . . . .	491
Viewing logical paths. . . . .	336	Viewing predefined reports. . . . .	492
Adding logical paths . . . . .	337	Creating custom reports . . . . .	493
Adding logical paths by using a CSV file . . . . .	337	Viewing and creating reports in the Cognos software . . . . .	494
Removing logical paths . . . . .	340	Searching the predefined reports for a property or performance metric . . . . .	494
Setting up data replication . . . . .	340	Scenarios . . . . .	495
Sessions . . . . .	340	Analyzing the volumes on the most active hypervisors . . . . .	495
Session types . . . . .	349	Monitoring daily the performance of volumes . . . . .	496
Session commands . . . . .	383	Investigating a degradation in the performance of a storage pool . . . . .	497
Site awareness . . . . .	412	Identifying the relationships between a storage system and other resources. . . . .	498
Preserve Mirror option . . . . .	413	Configuring . . . . .	498
Creating sessions and adding copy sets. . . . .	414	Setting options for reports . . . . .	498
Completing session administration tasks . . . . .	427	Scheduling a report to run with predefined options . . . . .	499
Using the Metro Mirror heartbeat . . . . .	428	Creating a copy of a predefined report . . . . .	500
Exporting copy set data . . . . .	430	Starting the Cognos Configuration GUI. . . . .	500
Importing copy set data . . . . .	431	Sending reports by email . . . . .	501
Exporting historical data for Global Mirror role pairs . . . . .	431	Running reports by using the <b>trcmd</b> command . . . . .	502
Modifying the location of session sites . . . . .	432	Specifying locations for saving reports . . . . .	503
Removing sessions . . . . .	433	Setting up multiple connections to the TPCDB data source . . . . .	505
Removing copy sets . . . . .	433	Customizing the logo and title of reports . . . . .	506
Migrating existing hardware relationships to Tivoli Storage Productivity Center for Replication . . . . .	434	Creating a search index of reports, report views, and folders . . . . .	507
Practicing disaster recovery. . . . .	434	Configuring security for the Cognos software with the default authentication configuration. . . . .	508
Practice volumes . . . . .	435		
Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session. . . . .	435		

Storage virtualizer reports that are available when you upgrade . . . . .	516	Predefined reports about storage tiering . . . . .	761
Storage virtualizer resource reports . . . . .	517	Running the Storage Resource Group - VDisk Workload Activity Details report . . . . .	762
Storage virtualizer performance reports. . . . .	530	Storage Resource Group - VDisk Workload Activity Details report . . . . .	763
Storage virtualizer performance details reports . . . . .	540	Running the VDisk Details report . . . . .	767
Predefined reports about historical information . . . . .	555	VDisk Details report . . . . .	767
Storage virtualizer custom reports about performance. . . . .	570	Running the MDisk Groups - VDisk Workload Activity report . . . . .	769
Predefined reports about resource relationships . . . . .	627	MDisk Groups - VDisk Workload Activity report . . . . .	770
Running the Storage Resource Relationships Summary report . . . . .	627	Running the Storage Virtualizer - VDisk Workload Activity report . . . . .	774
Storage Resource Relationships Summary report . . . . .	628	Storage Virtualizer - VDisk Workload Activity report . . . . .	774
Running the Storage Resource Relationships Summary (Configurable) report . . . . .	629	Running the Storage Resource Group - VDisk Workload Activity report . . . . .	778
Storage Resource Relationships Summary (Configurable) report . . . . .	630	Storage Resource Group - VDisk Workload Activity report . . . . .	778
Running the File System to Volume Relationships report . . . . .	632	Running the MDisk Group - VDisk Workload Activity Details report . . . . .	783
File System to Volume Relationships report . . . . .	632	MDisk Group - VDisk Workload Activity Details report . . . . .	783
Predefined reports about switches and switch ports . . . . .	634	Running the MDisk Group Details report . . . . .	787
Predefined reports about switches . . . . .	634	MDisk Group Details report . . . . .	787
Predefined reports about switch ports . . . . .	640	Running the MDisk Groups - Workload Activity report . . . . .	790
Predefined reports about groups . . . . .	648	MDisk Groups - Workload Activity report . . . . .	790
Running the Capacity of One Group report . . . . .	648	Read I/O capability formula . . . . .	795
Capacity of One Group report. . . . .	648	Changing threshold values in storage tier reports . . . . .	797
Running the Groups Capacity report . . . . .	650	Threshold values in storage tier reports . . . . .	798
Groups Capacity report . . . . .	650	Custom reports about performance . . . . .	798
Predefined reports about hypervisors . . . . .	651	Creating custom performance reports . . . . .	799
Running the Hypervisors Capacity report . . . . .	651	Data and properties in performance reports . . . . .	800
Hypervisors Capacity report . . . . .	652	Custom reports about capacity and relationships . . . . .	927
Running the Most Active Hypervisors report . . . . .	653	Creating custom capacity and relationship reports . . . . .	927
Most Active Hypervisors report . . . . .	654	Adding filters to reports. . . . .	929
Running the Summarized Performance of Volumes by Hypervisor report. . . . .	654	Data and properties in capacity and relationship reports . . . . .	933
Summarized Performance of Volumes by Hypervisor report . . . . .	655	Troubleshooting reports . . . . .	1007
Running the Performance of Volumes by Hypervisor report . . . . .	657	Showing the package version number and build ID. . . . .	1013
Performance of Volumes by Hypervisor report . . . . .	657	Reporting in a vSphere environment . . . . .	1013
Running the Hypervisor Data Stores Capacity report . . . . .	658	Viewing vSphere storage mapping to Tivoli Storage Productivity Center . . . . .	1014
Hypervisor Data Stores Capacity report . . . . .	659	Viewing storage system metrics in vSphere Web Client extension . . . . .	1017
Running the Hypervisor Disks Capacity report . . . . .	660	Viewing Tivoli Storage Productivity Center fabric connections in vSphere Web Client extension . . . . .	1024
Hypervisor Disks Capacity report . . . . .	660	Viewing block storage reports in the vSphere Web Client . . . . .	1025
Predefined reports about servers and file systems on servers . . . . .	661	Viewing file system reports in the vSphere Web Client . . . . .	1026
Predefined reports about servers . . . . .	661	Viewing Tivoli Storage Productivity Center VASA provider reports . . . . .	1027
Predefined reports about file systems on servers . . . . .	669	Reporting with the Tivoli Storage Productivity Center Monitoring Agent . . . . .	1028
Predefined reports about storage systems and components . . . . .	671		
Predefined reports about storage systems . . . . .	671		
Predefined reports about controllers, modules, and nodes . . . . .	685		
Predefined reports about disks . . . . .	696		
Predefined reports about host connections. . . . .	704		
Predefined reports about I/O groups . . . . .	710		
Predefined reports about managed disks . . . . .	715		
Predefined reports about storage pools . . . . .	723		
Predefined reports about ports . . . . .	737		
Predefined reports about RAID arrays . . . . .	743		
Predefined reports about volumes . . . . .	748		

<b>Appendix A. Reference</b>	<b>1035</b>
Alerts	1035
Triggering conditions for alerts	1036
Triggering actions and notifications for alerts	1064
Alert display colors	1066
Performance data	1069
Performance metrics	1069
Performance thresholds	1100
Cinder driver for Tivoli Storage Productivity Center	1126
Planning to deploy the Cinder driver	1127
Setting up the Cinder driver environment	1128
Requirements for installing the Cinder driver	1130
Cinder driver package contents	1132
Installing the Cinder driver	1132
Uninstalling the Cinder driver	1134
Configuring the Cinder driver	1134
Learn more about the Cinder driver	1135
Limitations and known issues for the Cinder driver	1136
Opening Tivoli Storage Productivity Center on Windows operating systems	1137
Opening Tivoli Storage Productivity Center CLIs and GUIs	1137
Accessing administration tools	1139
<b>Appendix B. Accessibility features for Tivoli Storage Productivity Center</b>	<b>1143</b>
<b>Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication</b>	<b>1145</b>
<b>Notices</b>	<b>1149</b>
Trademarks	1151

Privacy policy considerations	1151
-------------------------------	------

<b>Glossary</b>	<b>1153</b>
A	1153
B	1153
C	1153
D	1154
E	1154
F	1154
G	1154
H	1154
I	1155
J	1155
K	1155
L	1155
M	1155
N	1156
O	1156
P	1156
R	1156
S	1157
T	1157
V	1157
Z	1158

<b>Index</b>	<b>1159</b>
--------------	-------------





---

## About this guide

IBM® Tivoli® Storage Productivity Center manages storage infrastructure by centralizing, automating, and simplifying the management of complex and heterogeneous storage environments. IBM Tivoli Storage Productivity Center for Replication helps you manage data-copy functions. This guide provides task-oriented user information that helps you to obtain optimal product performance.

---

## Who should read this guide

This publication is intended for administrators and users who want to learn how to use Tivoli Storage Productivity Center to monitor and manage storage resources and infrastructure.

Administrators and users must be familiar, or become familiar, with the following concepts:

- Storage Area Network (SAN)
- Simple Network Management Protocol (SNMP)
- Storage provisioning, optimization, performance, tiers, and cloud configuration
- Tivoli Storage Productivity Center for Replication
- IBM DB2®
- Jazz™ for Service Management
- IBM Tivoli Common Reporting
- vSphere Web Client extension and VASA

---

## Publications

A number of publications are provided with Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. The following section describes how to access these publications online.

### Accessing publications online

Information about installing, configuring, using, upgrading, and uninstalling Tivoli Storage Productivity Center and related products is available online.

Use Table 1 on page x to view and download these publications. Translated documents are available for some products.

Table 1. Locations of publications for Tivoli Storage Productivity Center and related products

Product	Online location
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	<p><a href="http://www.ibm.com/support/knowledgecenter/SSNE44/">http://www.ibm.com/support/knowledgecenter/SSNE44/</a></p> <p>In the navigation pane, click <b>Tivoli Storage Productivity Center</b> for the appropriate release.</p> <p>To obtain PDF documents, in the navigation pane click <b>Tivoli Storage Productivity Center &gt; Printable documentation</b>.</p> <p>To view previous versions of the IBM Tivoli Storage Productivity Center Information Center, go to the Tivoli Storage Productivity Center page in Tivoli Documentation Central.</p>
IBM Tivoli Storage Productivity Center for Replication for System z®	<p><a href="http://www.ibm.com/support/knowledgecenter/SSU9FR/">http://www.ibm.com/support/knowledgecenter/SSU9FR/</a></p> <p>In the navigation pane, click <b>Tivoli Storage Productivity Center for Replication for System z</b> for the appropriate release.</p> <p>To obtain PDF documents, in the navigation pane click <b>Tivoli Storage Productivity Center for Replication for System z &gt; Printable documentation</b>.</p> <p>To view previous versions of the IBM Tivoli Storage Productivity Center for Replication for System z Information Center, go to the Tivoli Storage Productivity Center page in Tivoli Documentation Central.</p>
IBM WebSphere® Application Server	<a href="http://www.ibm.com/support/knowledgecenter/SSEQTP/mapfiles/product_welcome_was.html">http://www.ibm.com/support/knowledgecenter/SSEQTP/mapfiles/product_welcome_was.html</a>
Jazz for Service Management	<a href="http://www.ibm.com/support/knowledgecenter/SSEKCU/welcome">http://www.ibm.com/support/knowledgecenter/SSEKCU/welcome</a>
IBM System Storage® DS3000, IBM System Storage DS4000®, or IBM System Storage DS5000	<a href="http://www.ibm.com/support/entry/portal/">http://www.ibm.com/support/entry/portal/</a>
IBM System Storage DS6000™	<p><a href="http://www.ibm.com/support/knowledgecenter/">http://www.ibm.com/support/knowledgecenter/</a></p> <p>Go to the preceding website and expand the twisties as follows: <b>System Storage &gt; Disk systems &gt; Enterprise Storage Servers &gt; DS6000 series</b>.</p>
IBM System Storage DS8000®	<p><a href="http://www.ibm.com/support/knowledgecenter/">http://www.ibm.com/support/knowledgecenter/</a></p> <p>Go to the preceding website and go to <b>System Storage &gt; Disk systems &gt; Enterprise Storage Servers</b>. Select the DS8000 series system that you want.</p>
IBM System Storage DS® Open Application Programming Interface publications	<p><a href="http://www.ibm.com/support/docview.wss?uid=ssg1S7001160">http://www.ibm.com/support/docview.wss?uid=ssg1S7001160</a></p> <p>Use these publications for information about how to install, configure, and use the CIM agent.</p>
IBM System Storage SAN Volume Controller	<a href="http://www.ibm.com/support/knowledgecenter/STPVGU">http://www.ibm.com/support/knowledgecenter/STPVGU</a>

Table 1. Locations of publications for Tivoli Storage Productivity Center and related products (continued)

Product	Online location
IBM Storwize® V3500	<a href="http://www.ibm.com/support/knowledgecenter/STLM6B/landing/v3500_welcome.html">http://www.ibm.com/support/knowledgecenter/STLM6B/landing/v3500_welcome.html</a>
IBM Storwize V3700	<a href="http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html">http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html</a>
IBM Storwize V5000	<a href="http://www.ibm.com/support/knowledgecenter/STHGJ/landing/V5000_welcome.html">http://www.ibm.com/support/knowledgecenter/STHGJ/landing/V5000_welcome.html</a>
IBM Storwize V7000	<a href="http://www.ibm.com/support/knowledgecenter/ST3FR7/landing/V7000_welcome.html">http://www.ibm.com/support/knowledgecenter/ST3FR7/landing/V7000_welcome.html</a>
IBM Storwize V7000 Unified	<a href="http://www.ibm.com/support/knowledgecenter/ST5q4U/landing/v7000_unified_welcome.html">http://www.ibm.com/support/knowledgecenter/ST5q4U/landing/v7000_unified_welcome.html</a>
IBM Scale Out Network Attached Storage (IBM SONAS)	<a href="http://www.ibm.com/support/knowledgecenter/STAV45/com.ibm.sonas.doc/sonasWelcome.htm?cp=STAV45%2F0-0">http://www.ibm.com/support/knowledgecenter/STAV45/com.ibm.sonas.doc/sonasWelcome.htm?cp=STAV45%2F0-0</a>
IBM XIV® Storage System	<a href="http://www.ibm.com/support/knowledgecenter/STJTAG">http://www.ibm.com/support/knowledgecenter/STJTAG</a>
IBM General Parallel File System and IBM System x® GPFS™ Storage Server	<a href="http://www-01.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html">http://www-01.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html</a>
IBM DB2 Database for Linux UNIX and Windows	<a href="http://www.ibm.com/support/knowledgecenter/SSEPGG/welcome">http://www.ibm.com/support/knowledgecenter/SSEPGG/welcome</a>
IBM System Storage N series	<a href="http://www.ibm.com/systems/storage/network/index.html">http://www.ibm.com/systems/storage/network/index.html</a>  For more information about IBM System Storage N series, see <a href="http://www.ibm.com/systems/storage/network/software/">http://www.ibm.com/systems/storage/network/software/</a> .
IBM Systems Director	<a href="http://www.ibm.com/systems/software/director/index.html">http://www.ibm.com/systems/software/director/index.html</a>
VMware vSphere	<a href="http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html">http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html</a>
IBM Cognos® Business Intelligence	<a href="http://www.ibm.com/support/knowledgecenter/SSEP7J">http://www.ibm.com/support/knowledgecenter/SSEP7J</a>
IBM FlashSystem V840 Enterprise Performance Solution	<a href="http://www.ibm.com/systems/storage/flash/v840/overview.html">http://www.ibm.com/systems/storage/flash/v840/overview.html</a>

## IBM Redbooks

The IBM Redbooks® are publications about specialized topics.

You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page at <http://www.redbooks.ibm.com>.

## Translation

Translated publications are available from IBM Knowledge Center, which is available in certain translated languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a translated version, the information is displayed in English, which is the default language. Translations of the PDFs are available when the information is translated.

## Printable documentation


Print a single topic, using either the print function of your web browser or the print function of the IBM Tivoli Storage Productivity Center Information Center. To print multiple topics easily, print from the PDF version of this documentation.

### Viewing, printing, and saving a PDF

The topics in the IBM Tivoli Storage Productivity Center Information Center are captured in a series of PDF documents that are accessible from this topic. Printing from a PDF is the best way to print multiple topics at one time.

1. Download Adobe Acrobat Reader, if necessary.
2. View and print one of the available PDF files.
3. Save the PDF file to open it in Acrobat Reader without a browser.

### Downloading Adobe Acrobat Reader

To view or print these PDF files, you need Adobe Acrobat Reader. You can download a free copy from the Adobe website  ([www.adobe.com/products/acrobat/readstep.html](http://www.adobe.com/products/acrobat/readstep.html)).

### Viewing and Printing PDF files

Click one of the following links to view the corresponding PDF within your browser. Print the PDF by selecting the **Print** icon on the browser toolbar, or by selecting **File > Print**.

- IBM Tivoli Storage Productivity Center Installation Guide (about 5000 KB)
- IBM Tivoli Storage Productivity Center Quick Installation Guide (about 3000 KB)
- IBM Tivoli Storage Productivity Center User's Guide (about 12046 KB)
- IBM Tivoli Storage Productivity Center Administrator's Guide (about 5346 KB)
- IBM Tivoli Storage Productivity Center Problem Determination Guide (about 900 KB)
- IBM Tivoli Storage Productivity Center Quick Start Guide (about 90 KB)
- IBM Tivoli Storage Productivity Center Select Edition Quick Start Guide (about 90 KB)
- IBM SmartCloud Virtual Storage Center V5.2 Storage Analytics Engine Quick Start Guide (about 90 KB)

### Saving a PDF file

To save a PDF file, complete the following steps:

1. Right-click the link to the PDF in your browser.
2. Do one of the following tasks:

#### For Internet Explorer

Right-click the link and click **Save Target As**.


#### For Netscape Navigator or Mozilla Firefox

Right-click the link and click **Save Link As**.

3. Navigate to the directory in which you want to save the PDF file.

4. Click **Save**.

## Printing a single topic


The **Print Page** button () is located on the border of the IBM Tivoli Storage Productivity Center Information Center.

1. Display the topic that you want to print by selecting it from the table of contents or from the search results.
2. (Recommended) Click the displayed topic to place focus on the content area of the IBM Tivoli Storage Productivity Center Information Center.
3. Click the **Print Page** button.

## Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information.

To provide comments or suggestions about the product documentation, go to

 [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/). Click **Add Comment** on the documentation welcome page or the individual topic pages. You must have an IBM ID and password to add a comment.

---

## Tivoli Storage Productivity Center Service Management Connect community

Connect, learn, and share with Service Management professionals: product support technical experts who provide their perspectives and expertise.

Access Service Management Connect at <https://www.ibm.com/developerworks/servicemanagement/>. Use Service Management Connect in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the Storage Management community.
- Read blogs to benefit from the expertise and experience of others.
- Use wikis and forums to collaborate with the broader user community.

---

## Contacting IBM Software Support

You can contact IBM Software Support by phone, and you can register for support notifications at the technical support website.

- Go to the IBM Tivoli Storage Productivity Center technical support website at [http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli\\_Storage\\_Productivity\\_Center](http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center).

To receive future support notifications, sign in under **Subscribe to support notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for Tivoli Storage Productivity Center technical support website updates.

- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to the Tivoli Storage Productivity Center technical support website to find support by country. Expand **Contact support** and click **Directory of worldwide contacts**.

You can also review the *IBM Software Support Handbook*, which is available at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

The support website offers extensive information, including a guide to support services; frequently asked questions (FAQs); and documentation for all IBM Software products, including Redbooks and white papers. Translated documents are also available for some products.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information might also be needed to access various online services available on the website. See “Reporting a problem.”

## Reporting a problem

Provide the IBM Support Center with information about the problems that you report.

Have the following information ready when you report a problem:

- The IBM Tivoli Storage Productivity Center version, release, modification, and service level number.
- The communication protocol (for example, TCP/IP), version, and release number that you are using.
- The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.
- The exact text of any error messages.

---

## Conventions used in this guide

Information is given about the conventions that are used in this publication.

This publication uses several conventions for special terms and actions, and for operating system-dependent commands and paths.

The following typeface conventions are used in this publication:

### **Bold**

- Flags that display with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

### *Italic*

- Variables
- Values that you must provide
- New terms
- Words and phrases that are emphasized
- Titles of documents

### **monospace**

- Commands and command options
- Flags that display on a separate line
- Code examples and output
- Message text

- Names of files and directories
- Text strings that you must type, when they display within text
- Names of Oracle Java™ methods and classes
- HTML and XML tags that display like this, in monospace type

For syntax notations, remember the following details.

- In AIX®, the prompt for the root user is #.
- In AIX and Linux, the commands are case-sensitive, so you must type commands exactly as they are shown.





## New for Tivoli Storage Productivity Center Version 5.2.4

New features and enhancements are available in IBM Tivoli Storage Productivity Center Version 5.2.4.

### New features for Tivoli Storage Productivity Center

Learn about the following features and updates in Tivoli Storage Productivity Center Version 5.2.4. These features are available for all product licenses.

- Navigating the GUI
- Defining alerts for configuration changes and performance problems on monitored resources
- Monitoring and managing applications and departments
- Viewing information about GPFS snapshots, inodes, resources, and status changes
- Viewing information about multiple paths on server or hypervisor disks
- Currency

#### Navigating the GUI

You can use the menu bar to open resource pages and access the main product features. You can identify the ID of the user who is logged in to the GUI, view the product license that is active, and access the online help.

The following figure shows the menu bar with the **Storage** menu open:

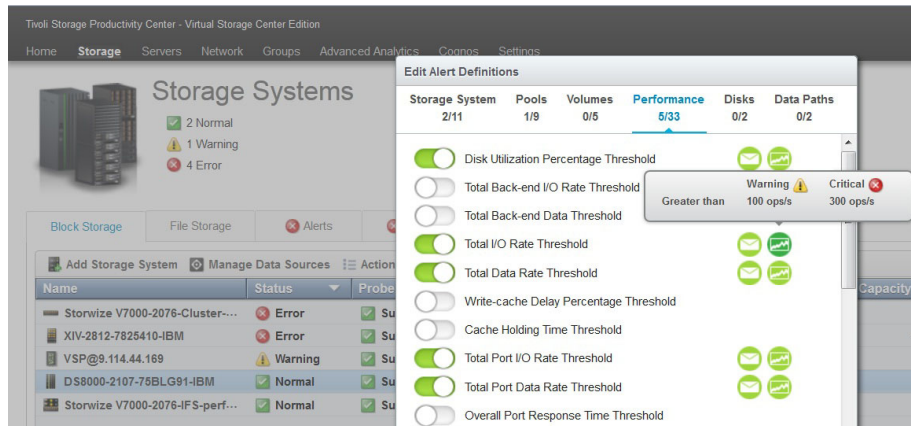


See [Navigation](#).

#### Defining alerts for configuration changes and performance problems on monitored resources

You can use the web-based GUI to determine when and how you are alerted to configuration changes or performance problems on monitored resources. Alerts are triggered automatically when conditions are detected during data collection and event processing.

You can view and edit the alert definitions for each resource in one, central location. For example, you can right-click a storage system and define a condition that triggers an alert when that storage system goes offline, and define a threshold that triggers an alert when the total I/O rate of its components exceeds 300 operations per second. You can also define how you are notified about each alert, or set a global notification setting for all of the alert definitions that are associated with a resource.



See [Migrating alerts and “Alerting”](#) on page 159.

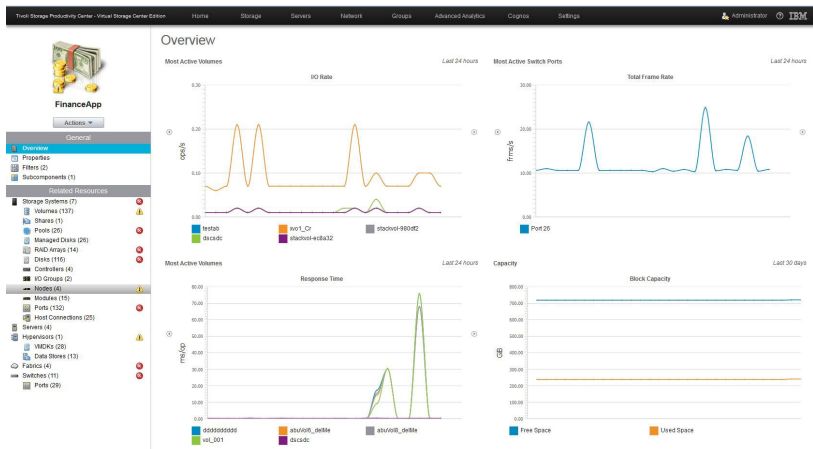
## Monitoring and managing applications and departments

You can use the web-based GUI to create applications and departments to model your business hierarchy. You can monitor the storage capacity that is being used by applications and departments and understand the SAN resources that your applications depend on in order to troubleshoot performance from the perspective of your applications.

### Modeling applications

You can create applications that model the storage usage of the applications within your business organization. You can then view capacity usage, performance, and SAN resource dependencies on a per application basis. You can associate applications with departments.

The following figure shows the new application Overview page, which includes a list of related components and storage resources, as well as capacity trending and performance charts for the monitored application:



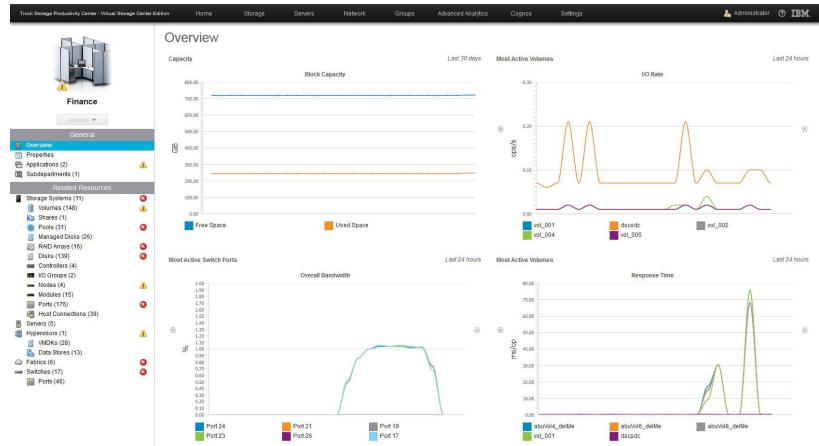
See “Applications” on page 225.

### Modeling departments

You can create departments and monitor the storage capacity that is used by the departments in your business environment. You can structure departments hierarchically to mirror your business structure. By adding departments to the hierarchy, you can regulate

and compare what departments have the most capacity usage and track the trending of the storage that is used by specific departments.

The following figure shows the new department Overview page, which includes a list of related components and storage resources, as well as capacity trending and performance charts for the monitored department:



See “Departments” on page 233.

## Viewing information about GPFS snapshots, inodes, resources, and status changes

You can view the following information about GPFS clusters:

### Viewing information about GPFS snapshots

You can view information about the snapshots that are associated with a GPFS cluster. For example, you can view the file system and fileset that were used to create the snapshot, the status of the snapshot, and the amount of space that is used by the snapshot.

### Viewing the inodes by file system chart

For storage systems that are configured for file storage, for example, Storwize V7000 Unified or GPFS clusters, you can view a chart that shows the inode utilization of file systems. Use this chart, on the storage system details page, to quickly identify the file systems that are running out of available inodes.

### Viewing resources that are related to GPFS clusters

You can view information about the following resources that are associated with a GPFS cluster:

- Volumes on back-end storage systems that provide storage to Network Shared Disks (NSDs) on a GPFS cluster
- Servers that are members of a GPFS cluster

### Viewing status changes that are detected on GPFS clusters

Changes to the GPFS cluster are detected within minutes. The following changes are detected and trigger alerts that you can view in the web-based GUI:

- Adding or removing a node or NSD
- A status change for a node or NSD
- A GPFS cluster goes offline or comes online

## Viewing information about multiple paths on server or hypervisor disks

You can view the following information about multiple paths on server or hypervisor disks:

- The aggregated status of multiple paths on server or hypervisor disks.
- The multipathing policy for a disk.
- All of the paths on a server or hypervisor disk, and whether a path is the preferred path for a disk.

You can also include multipathing properties and status information about server and hypervisor disks in capacity and relationship reports.

## Currency

Tivoli Storage Productivity Center includes the following new platform support:

- Expanded support for all Linux flavors that are supported by GPFS 4.1.
- XIV systems Gen3 11.5 with cloud-ready storage, data protection, high availability, and simplified management for multiple tenants.

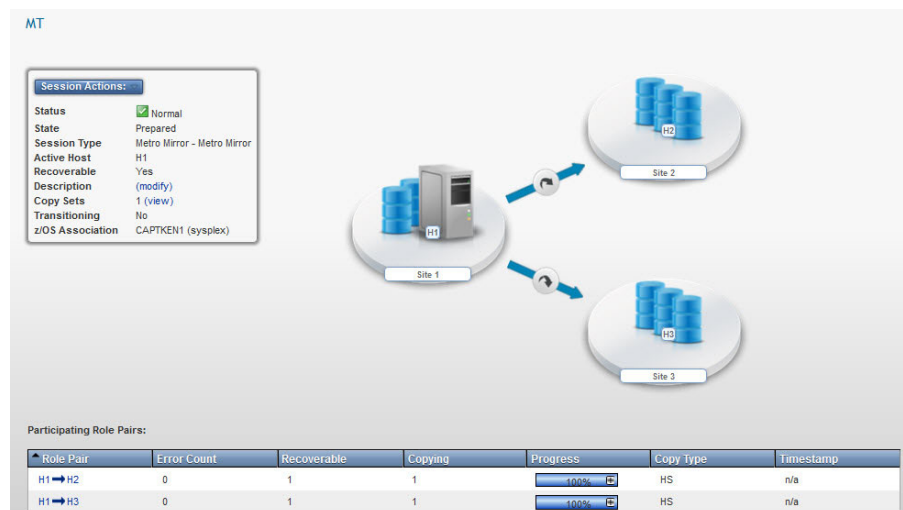
## New features for Tivoli Storage Productivity Center for Replication

Learn about the features and updates in Tivoli Storage Productivity Center for Replication V5.2.4.

### Creating multi-target sessions for System Storage DS8000 storage systems

A new Metro Mirror session type is available for IBM System Storage DS8000 storage systems that support single source to multi-target relationships. This session type is called Metro Mirror - Metro Mirror. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.

The following figure shows the Session Details page for a multi-target session:



See “Metro Mirror” on page 358, “Completing session administration tasks” on page 427, “Metro Mirror multi-target commands” on page 406, and “Metro Mirror - Metro Mirror session properties” on page 468.

### Defining port pairings for logical paths

You can use the `portpairings.csv` file to define the following types of port pairings:

- Port pairings for a logical subsystem to logical subsystem mapping.
- Port pairings for all logical subsystems that are in a session.
- Port pairings in which the pairings for all even logical subsystems and all odd logical subsystems are different.

The sample CSV file that is provided with Tivoli Storage Productivity Center for Replication, `portpairings.csv.sample`, includes examples of each type of port pairing.

See “Adding logical paths by using a CSV file” on page 337.

## **Limitations and known issues for Tivoli Storage Productivity Center**

For information about limitations and known issues that you might encounter when you use IBM Tivoli Storage Productivity Center Version 5.2.4, see the <http://www.ibm.com/support/docview.wss?uid=swg21684425> document.

## **Product licenses for Tivoli Storage Productivity Center**

Tivoli Storage Productivity Center offers product licenses with differing functionality to meet your storage management needs. For information about these licenses, see Product licenses.

### **Related information:**

[New for Tivoli Storage Productivity Center Version 5.2.3](#)

[New for Tivoli Storage Productivity Center Version 5.2.2](#)

[New for Tivoli Storage Productivity Center Version 5.2.1](#)

[New for Tivoli Storage Productivity Center Version 5.2](#)



---

## Chapter 1. Getting started

Learn about the key tasks for setting up Tivoli Storage Productivity Center for managing the resources and infrastructure in a storage environment. The getting started guide provides a list of tasks in sequential order for helping you to learn how to install, configure, and use Tivoli Storage Productivity Center.

Complete the following tasks to install and set up Tivoli Storage Productivity Center for monitoring and managing storage resources. Click the included links to learn more information about each task.

Table 2. Tasks for getting started with Tivoli Storage Productivity Center








Task	Description	 Learn more
Planning for installation and upgrades.	<p>The time that you invest in planning directly affects your ease of implementation. Planning includes the following tasks:</p> <ul style="list-style-type: none"><li>• Ensuring that your systems meet hardware and software requirements</li><li>• Becoming familiar with the operating systems, storage resources, agents, and communication protocols that can be used with Tivoli Storage Productivity Center</li><li>• Ensuring that ports are open for use</li><li>• Understanding authentication and role-based authorization</li><li>• Understanding the requirements for specific types of resources</li></ul>	<ul style="list-style-type: none"><li>•  Hardware requirements</li><li>•  Software requirements.</li><li>•  TCP/IP ports used by Tivoli Storage Productivity Center</li><li>•  Planning for Tivoli Storage Productivity Center authentication and authorization</li><li>•  Planning for resources</li></ul> <p>For more information about the software and resources that are supported by Tivoli Storage Productivity Center, see  Tivoli Storage Productivity Center Interoperability Matrix.</p>

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)







Task	Description	 Learn more
Installing or upgrading Tivoli Storage Productivity Center.	<p>If you have a previous version of Tivoli Storage Productivity Center installed, you can upgrade it to the current version. If you do not have a previous version, you can use the following methods to install it:</p> <p><b>Single-server environment</b> Install Tivoli Storage Productivity Center and its components on a single server.</p> <p><b>Multiple-server environment</b> Install Tivoli Storage Productivity Center on multiple servers if you are monitoring large storage environments, where one server is not sufficient to manage all the Tivoli Storage Productivity Center components.</p> <p>When you install Tivoli Storage Productivity Center, the following components are also installed:</p> <p><b>DB2 (required)</b> Tivoli Storage Productivity Center stores the information that it collects about storage resources in a DB2 database.</p> <p><b>Jazz for Service Management and Tivoli Common Reporting (optional)</b> Install these components if you want to use Tivoli Storage Productivity Center reports in the web-based GUI.</p>	<p>You can run the installation process as a GUI wizard or through a command line in silent mode.</p> <ul style="list-style-type: none"> <li>•  <a href="#">Upgrading and migrating Tivoli Storage Productivity Center</a></li> <li>•  <a href="#">Installing DB2</a></li> <li>•  <a href="#">Installing Tivoli Storage Productivity Center in a single-server environment</a></li> <li>•  <a href="#">Installing Tivoli Storage Productivity Center in a multiple-server environment</a></li> <li>•  <a href="#">Install Jazz for Service Management and Tivoli Common Reporting</a></li> </ul>



Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)


Task	Description	 Learn more
Accessing the Tivoli Storage Productivity Center GUIs.	<p>Tivoli Storage Productivity Center provides the following GUIs for managing and monitoring storage resources:</p> <p><b>Web-based GUI</b>            Contains many of the functions for monitoring, troubleshooting, and reporting on the condition of storage resources. This interface runs in a web browser and can be accessed from anywhere that you have a web browser and connectivity to a network.</p> <p><b>Stand-alone GUI</b>            The stand-alone GUI is the same, Oracle Java based interface that was available in Tivoli Storage Productivity Center v4.2.2.3 and earlier. In version 5.2 and later, many of the functions for monitoring and managing storage in this interface were moved to the web-based GUI.</p> <p><b>Tivoli Storage Productivity Center for Replication interfaces</b>            Tivoli Storage Productivity Center for Replication provides a GUI that runs in a web browser and a command-line interface. Use these interfaces to manage data replication and disaster recovery.</p>	<ul style="list-style-type: none"> <li>• “Starting Tivoli Storage Productivity Center” on page 42</li> <li>• “Available functions in the interfaces” on page 34</li> </ul>
Adding resources for monitoring.	<p>Before you can monitor the resources in your environment, you must add them to Tivoli Storage Productivity Center.</p> <p><b>Web-based GUI</b>            In the web-based GUI, you can add the following resources for monitoring:</p> <ul style="list-style-type: none"> <li>• Storage systems</li> <li>• Servers</li> <li>• Hypervisors</li> <li>• Switches</li> <li>• Fabrics</li> </ul> <p><b>Stand-alone GUI</b>            In the stand-alone GUI, you can add the following resources for monitoring:</p> <ul style="list-style-type: none"> <li>• NAS filers</li> <li>• Tivoli Storage Productivity Center servers as subordinate servers</li> </ul>	<p>“Adding resources” on page 127</p>

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)




Task	Description	 <b>Learn more</b>
Deploying Storage Resource agents on servers (optional).	In the web-based GUI, you can deploy Storage Resource agents on servers that you want to monitor. Deploy agents if you want to complete the following tasks: <ul style="list-style-type: none"> <li>• Collect storage usage and trending information about servers (scans)</li> <li>• Collect database application information</li> <li>• Discover the fabrics that are visible to a server</li> <li>• Enforce storage policies</li> <li>• Collect zone information</li> </ul>	For more information about Storage Resource agents, see “Adding servers” on page 138.  For a list of the servers where you can deploy Storage Resource agents, see  Tivoli Storage Productivity Center Interoperability Matrix and go to the <i>Agents, Servers, and Browsers</i> section.
Collecting data about resources.	Determining the information that you want to gather about storage resources is critical to helping you implement a storage management strategy. When you determine the types of information you want to collect about your resources, you can select the data collection jobs that you want to run. <b>Tip:</b> In the web-based GUI, probes and performance monitors are automatically scheduled when you add resources for monitoring:  <b>Probes</b> Probes collect status, asset, and capacity information about resources.  <b>Performance Monitors</b> Performance monitors collect metrics that measure the performance of storage systems and switches.  In the stand-alone GUI, you must manually schedule pings, scans, and TPC server probes after you add resources:  <b>Scans</b> Scans collect file usage and trending information about servers or hypervisors.  <b>Pings</b> Pings collect information about the availability (uptime) of servers.  <b>TPC Server probes</b> TPC server probes gather the information that was collected by multiple Tivoli Storage Productivity Center servers.	<ul style="list-style-type: none"> <li>• “Adding resources” on page 127</li> <li>• “Collecting data” on page 145</li> <li>•  Collecting storage statistics from Tivoli Storage Productivity Center servers</li> </ul>

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)






Task	Description	 Learn more
Configuring how alerts are received.	In the web-based GUI, enable alert notifications. Alerts can define notification actions that send email, generate Simple Network Management Protocol (SNMP) traps, or generate IBM Tivoli Netcool/OMNIBus events. To receive these types of alert notifications, you must configure Tivoli Storage Productivity Center to direct the notifications to a specific email address, SNMP destination, or Tivoli Netcool/OMNIBus server.	 <a href="#">Configuring alert notifications</a>
Specifying the alerts that are generated when a problem is detected on monitored storage.	<p>Determining when and how you are alerted to conditions or violations within your storage environment is important to helping you maintain and administer storage resources. Alerts are triggered by conditions that are detected during data collection and event processing.</p> <p>To set up and manage alerts, complete the following tasks:</p> <ul style="list-style-type: none"> <li>• For each resource that you want to monitor for issues, specify the conditions and performance thresholds that can trigger alerts.</li> <li>• Determine how you want to be notified when alerts are triggered.</li> <li>• In the web-based GUI, view, acknowledge, and remove alerts that are detected. For alerts that are triggered by performance threshold violations, you can view a chart of the historical performance of the resource that violated a threshold.</li> </ul>	<ul style="list-style-type: none"> <li>• “Defining alerts” on page 162</li> <li>• “Viewing and modifying alert notification settings” on page 174</li> <li>• “Viewing all alerts” on page 181</li> <li>• “Triggering conditions for alerts” on page 1036</li> </ul>
Configuring data retention (optional).	In the web-based GUI, specify how long to retain the data that is collected about resources and the log files that are generated by Tivoli Storage Productivity Center. By specifying the number of weeks for history retention, you can control the amount of data that is retained and available for historical analysis and charting. The longer that you retain data, the more informative your analysis, but the more storage space that is required to store that data. Default retention settings are provided.	 <a href="#">Configuring history and data retention</a>

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)

Task	Description	 Learn more
Configuring authentication and user access.	Determine the authentication mechanism and authorization level for users. The authentication mechanism determines the user groups to which you can assign roles. Roles determine the Tivoli Storage Productivity Center functions that are available to the users in a group. In the web-based GUI, you can assign the following roles to groups of users: <ul style="list-style-type: none"> <li>• Administrator</li> <li>• Monitor</li> <li>• External Application</li> </ul>	 Authorizing users

## Managing resources

Learn about the key tasks for getting started with monitoring and managing storage resources.

After installing Tivoli Storage Productivity Center and adding resources for monitoring, complete the following tasks to manage your storage environment.

Table 3. Tasks for monitoring and managing resources



Task	Description	 Learn more
Monitoring the performance of resources.	Collect and view key performance metrics about resources and specify alerts to be notified of threshold violations. Use metrics to measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.	<ul style="list-style-type: none"> <li>• “Monitoring the performance of resources” on page 205</li> <li>• Chapter 7, “Reporting,” on page 483</li> <li>• “Tracking performance issues on a storage system” on page 110</li> </ul>
Monitoring the capacity of resources.	View detailed information about the capacity of monitored resources in the web-based GUI and its reporting function. You can view capacity information is available for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual machines. <p>For example, you can view the total amount of storage space that is committed to a volume, the amount of allocated space that is used by a volume, and the amount of space that is allocated to a volume and is not yet used. In the reporting function of the web-based GUI, you can also view and create capacity reports about multiple resources.</p> <p>Use the information that is collected about storage capacity to help with the capacity planning of the storage within your organization.</p>	<ul style="list-style-type: none"> <li>• “Viewing information about resources” on page 185</li> <li>• Chapter 7, “Reporting,” on page 483</li> <li>•  Click the Help icon in the web-based GUI to view descriptions of the capacity information that is displayed for resources.</li> </ul>

Table 3. Tasks for monitoring and managing resources (continued)



Task	Description	 Learn more
Monitoring the status of resources.	<p>Monitor the status and condition of resources to identify potential problem areas in a storage environment. On the dashboard in the web-based GUI, you can view a high-level summary of the statuses for each type of resource. From the dashboard, you can drill down to view the individual status for each monitored resource.</p>	<ul style="list-style-type: none"> <li>• “How status is determined” on page 197</li> <li>• Viewing the overall status of a storage environment</li> <li>• Viewing the status of individual resources</li> <li>• Viewing the status of internal resources</li> </ul>
Viewing the relationships between resources.	<p>In the web-based GUI, you can generate a graphical or tabular representation of the relationships between host connections, servers, hypervisors, virtual machines, storage systems, and the fabrics through which they are connected. View data paths to monitor status and pinpoint problem areas in the path of monitored storage throughout your environment.</p> <p>In the reporting tool of the web-based GUI, view reports about the relationships between resources. For example, you can view a report about the end-to-end relationship of all the servers in your environment to back-end storage.</p> <p>In the stand-alone GUI, use the topology viewer to view the relationships between resources.</p>	<ul style="list-style-type: none"> <li>• Viewing data paths in the web-based GUI</li> <li>• Viewing relationship reports in the web-based GUI</li> <li>•  Using the topology viewer</li> </ul>
Monitoring and managing applications and departments.	<p>Storage resources can be grouped together to represent data modeling concepts such as applications, subcomponents (member applications), departments and subdepartments to create a business hierarchy. The monitoring and management of applications and departments enables you to perform capacity trending, health monitoring, and performance troubleshooting for the storage resources in your environment.</p> <p>In the web-based GUI, you can create applications and departments, add applications and departments, use a filtering criteria to add resources to the applications, remove related resources from within the context of applications, remove applications and departments, view key properties and view detailed storage information for applications and departments.</p>	<ul style="list-style-type: none"> <li>• Creating applications</li> <li>• Creating subcomponents</li> <li>• Adding applications to departments</li> <li>• Adding resources to applications with filters</li> <li>• Removing resources from applications</li> <li>• Viewing information about applications</li> <li>• Creating departments</li> <li>• Adding subdepartments</li> <li>• Viewing information about departments</li> </ul>

Table 3. Tasks for monitoring and managing resources (continued)





Task	Description	 Learn more
Identifying hot spots of high utilization.	<p>Tivoli Storage Productivity Center collects detailed information about storage system pools, including utilization values. The utilization value is an estimate of the average daily utilization of the physical resources such as controllers, disks, ports, and nodes that are associated with the pool.</p> <p>By scanning the utilization values for pools on a storage virtualizer, such as SAN Volume Controller, you can identify the hot spots or the pools that have high utilization values on the storage virtualizer. You can also determine the course of action that you take.</p>	<ul style="list-style-type: none"> <li>• “Redistributing volume workloads to balance pools” on page 94</li> <li>• “Re-tiering volumes to balance pools” on page 95</li> <li>• “Tiering volumes by I/O density and file usage” on page 97</li> <li>• “Collocating volumes” on page 101</li> </ul>
Setting up self-service provisioning.	<p>Storage provisioning is the process of assigning storage to optimize the performance of a storage environment. To provision storage with Tivoli Storage Productivity Center, create service classes and capacity pools to classify and separate storage resources. After configuration, use the web-based GUI, vSphere Web Client extension, or other tool with those service classes and capacity pools to provision storage.</p> <p>In the web-based GUI, you are guided through the steps for provisioning volumes or shares. Tivoli Storage Productivity Center then recommends storage placement that is based on storage system free space and performance data. If you agree with the recommendation, Tivoli Storage Productivity Center can implement the provisioning.</p>	<p>“Provisioning storage” on page 75</p>
Provisioning storage in a virtual storage environment.	<p>If you maintain virtual storage in a VMware vSphere environment, you can use the vSphere Web Client extension for Tivoli Storage Productivity Center to complete provisioning tasks and to view information about storage resources in the vSphere Web Client. The client includes details about performance and asset, copy services, and fabric information.</p>	<ul style="list-style-type: none"> <li>• “Provisioning storage with the vSphere Web Client extension” on page 266</li> <li>• “Limiting a set of storage to provision in the vSphere Web Client extension” on page 112</li> </ul>
Managing replication.	<p>Use Tivoli Storage Productivity Center for Replication to manage the advanced copy services for storage system through the following tasks:</p> <ul style="list-style-type: none"> <li>• Setting up management servers</li> <li>• Adding storage connections and paths</li> <li>• Creating sessions and copy sets</li> <li>• Setting up data replication</li> <li>• Performing planned maintenance</li> <li>• Recovering from a disaster</li> </ul>	<ul style="list-style-type: none"> <li>• Chapter 6, “Managing replication,” on page 303</li> <li>• “Replication scenarios” on page 114</li> </ul>

Table 3. Tasks for monitoring and managing resources (continued)

Task	Description	 Learn more
Determining the storage policies that you want to enforce.	<p>Enforcing storage policies within your organization can be critical to ensuring the most cost-effective use of your resources. In the stand-alone GUI, use quotas to define limits on the amount of storage that a user or a group of users can consume. Use constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.</p> <p>You must use the web-based GUI to deploy Storage Resource agents on the servers where you want to enforce policies.</p>	 Creating and enforcing storage policies
Running reports and setting up report scheduling and delivery.	<p>Use reports to view detailed information about the condition, capacity, and performance of resources. Tivoli Storage Productivity Center provides multiple user interfaces for running reports.</p> <p><b>Web-based GUI</b>            Access the reporting function in the web-based GUI to run, schedule, and deliver predefined and custom reports. These reports include information about storage capacity, assets, and performance of storage systems, switches, fabrics, servers, hypervisors, and their internal resources.</p> <p><b>vSphere Web Client</b>            Use the vSphere Web Client extension to view reports about virtual environments and storage resources. You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems.</p> <p><b>Stand-alone GUI</b>            View system, batch, and online reports. These reports are organized into different types and categories and provide both summary and detailed information, depending on your needs.</p>	<ul style="list-style-type: none"> <li>• Reporting in the web-based GUI</li> <li>• “Reporting in a vSphere environment” on page 1013</li> <li>•  Reporting in the stand-alone GUI</li> </ul>

**Related concepts:**

Chapter 5, “Managing resources,” on page 127












Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

## Configuring and administering

Learn about the key tasks for getting started with configuring and administering Tivoli Storage Productivity Center and its related components.

Complete any of the following tasks to customize Tivoli Storage Productivity Center for your environment and administer its components.

Table 4. Tasks for configuring and administering Tivoli Storage Productivity Center

Task	Description	 Learn more
Configuring Tivoli Storage Productivity Center.	<p>Configure Tivoli Storage Productivity Center according to the standards within your organization, as needed. The configuring tasks that you want to complete determines which GUI to use:</p> <p><b>Web-based GUI</b></p> <ul style="list-style-type: none"> <li>• Customizing lists of resources, tasks, and alerts</li> <li>• Enabling reports in Tivoli Common Reporting</li> <li>• Configuring how alerts are received</li> <li>• Configuring data retention</li> <li>• Configuring authentication and user access</li> </ul> <p><b>Stand-alone GUI</b></p> <ul style="list-style-type: none"> <li>• Defining rules for generating email addresses of quota violators</li> <li>• Configuring NAS filers for monitoring</li> <li>• Analyzing historical data in a SAN configuration</li> </ul> <p><b>vSphere client</b> Configure the vSphere client and VASA provider</p>	<ul style="list-style-type: none"> <li>• “Customizing lists of resources, tasks, and alerts” on page 287</li> <li>• If the reporting function is not enabled in the web-based GUI, you must install Jazz for Service Management and IBM Tivoli Common Reporting. For more information, see  <a href="#">Install Jazz for Service Management and Tivoli Common Reporting</a>.</li> <li>•  <a href="#">Configuration tasks in the stand-alone GUI</a></li> <li>•  <a href="#">Configuring the vSphere Web Client extension for Tivoli Storage Productivity Center</a></li> <li>•  <a href="#">Configuring the VASA provider</a></li> </ul>
Administering DB2 and Tivoli Storage Productivity Center.	<p>After Tivoli Storage Productivity Center is installed and configured, complete the following administering tasks as needed:</p> <ul style="list-style-type: none"> <li>• Administering resources and data sources</li> <li>• Starting and stopping Tivoli Storage Productivity Center servers</li> <li>• Checking the status of Tivoli Storage Productivity Center components</li> <li>• Packaging log files for IBM Software Support</li> <li>• Changing passwords for Tivoli Storage Productivity Center components</li> <li>• Administering DB2 and backing up the database repository</li> </ul>	<ul style="list-style-type: none"> <li>•  <a href="#">Administering resources and data sources</a></li> <li>•  <a href="#">Starting and stopping the Tivoli Storage Productivity Center servers</a></li> <li>•  <a href="#">Checking Tivoli Storage Productivity Center status</a></li> <li>•  <a href="#">Packaging log files for IBM Software Support</a></li> <li>•  <a href="#">Changing passwords</a></li> <li>•  <a href="#">Administering DB2</a></li> </ul>



---

## Navigation

The web-based GUI provides many of the functions for managing a storage environment. To access the functions in this GUI, use the menu bar at the top of the main window.

GUI element	Description
Page banner	<p>Use the banner at the top of every page in the GUI to complete the following tasks:</p> <ul style="list-style-type: none"><li>• View the name of the product license that is active.</li><li>• Access resource pages and main product functions from the menu bar.</li><li>• View the ID and role of the user who is logged in to the GUI. The role of a user determine the product functions that are available to that user. For more information about roles, see “Role-based authorization” on page 83.</li><li>• Access the online help to view information about the currently displayed page and the overall product.</li></ul>

GUI element	Description
Home	<p><b>Dashboard</b> Use the dashboard to view an overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about the status and capacity of resources, the jobs that collect performance data, and the unacknowledged alert conditions that are detected on resources.</p> <p><b>Performance Monitors</b> Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations.</p> <p><b>Alerts</b> Use alerts to be notified when certain conditions or violations are detected on monitored resources. In the stand-alone GUI, define alerts and thresholds. In the web-based GUI, view the alerts that were detected and the thresholds that were violated.</p> <p><b>Tasks</b> Manage the tasks that are created when you complete the following actions:</p> <ul style="list-style-type: none"> <li>• Assign storage to servers, hypervisors, or a cluster.</li> <li>• Place volumes on the tiers that match the workload requirements of the volume.</li> <li>• Move or convert volumes.</li> <li>• Move volumes from overutilized pools to pools that are less used.</li> </ul> <p><b>System Management</b> View information about the overall condition of Tivoli Storage Productivity Center. You can view information about the servers on which the product is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.</p>

GUI element	Description
Storage	<p>Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems:</p> <ul style="list-style-type: none"> <li>• Modify the schedules for data collection.</li> <li>• View the data path of connected resources.</li> <li>• Assign storage to servers, hypervisors, or a cluster.</li> <li>• Place volumes on the tiers that match the workload requirements of the volume.</li> <li>• Move or convert volumes.</li> <li>• Move volumes from overutilized pools to pools that are less used.</li> <li>• Add resources to a capacity pool.</li> <li>• Open management GUIs.</li> <li>• Update login credentials.</li> <li>• Acknowledge statuses.</li> </ul>
Servers	<p>Add servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> <li>• Deploy Storage Resource agents to servers for full monitoring.</li> <li>• Modify the schedules for data collection.</li> <li>• View the data path of connected resources.</li> <li>• Assign storage to servers, hypervisors, or a cluster.</li> <li>• Modify and upgrade Storage Resource agents.</li> <li>• Acknowledge statuses.</li> </ul>
Network	<p>Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> <li>• Modify the schedules for data collection.</li> <li>• View the data path of connected resources.</li> <li>• Setting or modifying the zoning policy.</li> <li>• Open the management GUI for a switch.</li> <li>• Acknowledge statuses.</li> </ul>

GUI element	Description
Groups	<p>Add applications and departments to combine resources that are grouped by applications and departments. Applications and departments are used to monitor and view assets, status, capacity and performance information and view information about their related resources. You can also complete the following actions to manage and administer applications and departments:</p> <ul style="list-style-type: none"> <li>• Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.</li> <li>• View the status of resources that make up the applications and departments.</li> <li>• Create resource filters to automatically add resources to applications.</li> <li>• Add resources directly to applications and add applications to existing applications and departments.</li> <li>• Add departments to existing departments.</li> <li>• View subcomponent (member application) information to support hierarchal levels of storage grouping.</li> </ul>
Advanced Analytics	<p>Configure Tivoli Storage Productivity Center for provisioning and optimizing storage. Configuration tasks include creating service classes and capacity pools to classify and separate resources. The web-based GUI includes the following wizards to guide you through the steps of managing your storage:</p> <ul style="list-style-type: none"> <li>• Use the Provision Storage wizard to assign storage to servers, hypervisors, or a cluster.</li> <li>• Use the Analyze Tiering wizard to manage storage resources efficiently and reduce storage costs.</li> <li>• Use the Analyze Tiering, Balance Pools, and Transform Storage wizards to optimize and transform storage volumes.</li> </ul>
Cognos	<p>Create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors in the Cognos reporting software. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor.</p> <p>Use the Cognos software to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports.</p> <p>Use the web-based GUI when you want to individually view details about resources, for example when you want to troubleshoot a resource.</p> <p>When you click a reporting option, a separate browser tab is opened.</p> <p>For information about how to get started with the reporting software, see Getting started with reporting.</p>

GUI element	Description
Settings	Complete the following actions to customize Tivoli Storage Productivity Center for your environment: <ul style="list-style-type: none"> <li>• Modify the authentication repository and assign Tivoli Storage Productivity Center roles to user groups. Roles determine the product functions that are available to users.</li> <li>• Specify how long to retain the data that is collected about resources and the log files that are generated by Tivoli Storage Productivity Center.</li> </ul>

**Navigating in the stand-alone GUI:** The stand-alone GUI provides a different set of navigation tools for accessing its functions. For more information, see

 Overview of the stand-alone GUI.

## Keyboard navigation in the web-based GUI

Most of the features of the web-based GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the web-based GUI:

### For navigating the web-based GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

### For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

### For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

### For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

### For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

**Restriction:** For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

---

## Chapter 2. Product licenses

Product licenses are available for IBM Tivoli Storage Productivity Center. Each license determines the functions that are accessible in the stand-alone GUI, web-based GUI, command-line interface, and the API. You can review the functions that are included with each license and determine which license meets your storage management needs.

**Restriction:** When you upgrade your license, all of the servers, except the web server are recycled. After you upgrade your license, to access the functionality in the license, you must stop and restart the web server.

### **IBM Tivoli Storage Productivity Center**

This license contains the following functions:

#### **Data management**

A file and capacity management solution for heterogeneous storage environments. Data management includes enterprise-wide reporting and monitoring, and automated capacity provisioning for Direct Attached Storage (DAS), network-attached storage (NAS), and Storage Area Network (SAN) environments.

#### **Disk management**

Disk management helps you manage SANs and heterogeneous storage from one console.

#### **Fabric management**

Fabric management helps you manage the SAN fabric that connects the host systems and applications to the storage devices. This feature is a comprehensive management solution for multi-vendor SANs and includes automatic resource and topology discovery, monitoring, and alerts.

#### **Tivoli Storage Productivity Center for Replication**

This license includes all the Tivoli Storage Productivity Center for Replication functions.

#### **Performance management**

Disk and Fabric performance management functions.

#### **Device support**

You can see a complete list of devices that can be used with Tivoli Storage Productivity Center at <https://www.ibm.com/support/docview.wss?uid=swg21386446>. Click the appropriate release under **Storage**.

### **IBM Tivoli Storage Productivity Center Select Edition**

This license contains all the functions of the IBM Tivoli Storage Productivity Center license. This license is priced by enclosure rather than by capacity.

### **IBM SmartCloud® Virtual Storage Center Storage Analytics Engine**

This license contains everything that is in the IBM Tivoli Storage Productivity Center license, and also includes advanced analytical functions. This license was previously called IBM Tivoli Storage Productivity Center Advanced. The user interfaces and messages still refer to this license as IBM Tivoli Storage Productivity Center Advanced.

In the web-based GUI, the following advanced analytical functions are available:

- Optimize storage tiering by using the Analyze Tiering wizard
- Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard
- Provision block storage by using the Provision Storage wizard.
- Enable automatic zoning to create zones during block storage provisioning to connect a server to a storage system

In the stand-alone GUI, the following advanced analytical functions are available:

- Configuration analysis
- Configuration history
- Scans
- Policy management
- Server probes
- Scans
- Profiles
- Data Manager for databases
- Data Manager for Chargeback

This license is only available as part of the IBM SmartCloud Virtual Storage Center. The complete IBM SmartCloud Virtual Storage Center solution bundles Tivoli Storage Productivity Center with IBM System Storage SAN Volume Controller and Tivoli Storage FlashCopy<sup>®</sup> Manager. In the IBM SmartCloud Virtual Storage Center solution, System Storage SAN Volume Controller provides a virtualization platform and remote replication functions, and the Tivoli Storage FlashCopy Manager provides data backup and restore capabilities. For more information about IBM SmartCloud Virtual Storage Center, see <http://www.ibm.com/software/products/us/en/vsc/>.

If you purchased a license for IBM Tivoli Storage Productivity Center Version 4.2, when you upgrade to Version 5.2, you can upgrade to the appropriate license. Table 5 shows the parallel functions in Version 5.2.

*Table 5. Tivoli Storage Productivity Center licenses when you upgrade from Version 4.2 to Version 5.2*

<b>If you have this Tivoli Storage Productivity Center Version 4.2 license</b>	<b>You are entitled to this Tivoli Storage Productivity Center Version 5.2 license</b>
Basic Edition	When you upgrade to the Version 5.2, you get the same functionality that you had with the Version 4 Basic Edition license.
Disk Select	IBM Tivoli Storage Productivity Center Select Edition
Disk	IBM Tivoli Storage Productivity Center
Replication 2-Site or Replication 3-Site	IBM Tivoli Storage Productivity Center
Data	No entitlement migration to Version 5.2. Contact your IBM sales representative for options.



Table 5. Tivoli Storage Productivity Center licenses when you upgrade from Version 4.2 to Version 5.2 (continued)

If you have this Tivoli Storage Productivity Center Version 4.2 license	You are entitled to this Tivoli Storage Productivity Center Version 5.2 license
Standard Edition	No entitlement migration to Version 5.2. Contact your IBM sales representative for options.

## Functions available in the web-based GUI based on role and license

When you use the web-based GUI, your Tivoli Storage Productivity Center role and product license determine the functions that are available.

Users who are assigned the Administrator role or the Monitor role can use the web-based GUI. The functions that are available depend on the role that is assigned to the user:

### Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

### Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs. Users who are assigned the Monitor role do not have access to administrative functions, with the following exceptions:

- They can assign storage pools to tiers.
- They can be granted permission in a service class to provision storage by using the service class. If so, they can use the Provision Storage wizard to create a provisioning task. Users can delete provisioning tasks that they create. If the service class specifies that administrator approval is not required, the users can run or schedule the provisioning tasks that they create.

To use some functions of the web-based GUI, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. If you have the Tivoli Storage Productivity Center license or the Tivoli Storage Productivity Center Select license, the following functions are not available:

- Storage tier optimization. The Analyze Tiers wizard is not available.
- Pool balancing. The Balance Pools wizard is not available.
- Block storage provisioning. You are not able to provision volumes by using either the Provision Storage wizard or the vSphere Web Client extension. However, there are no license restrictions for file storage provisioning, therefore, Network Attached Storage (NAS) file shares can be provisioned.

The following table outlines the functions available by role. Some functions are restricted to the Administrator role. Functions that are restricted to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license are also shown. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 6. Role and license restrictions in the web-based GUI

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Performance monitors	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about the switch or storage system that is being monitored</li> <li>• Open performance monitor logs</li> <li>• Export performance monitor information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Start or stop a performance monitor</li> <li>• Schedule a performance monitor</li> </ul>	All functions are available, depending on role.
Alerts	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about alerts</li> <li>• Acknowledge alerts</li> <li>• Export alert information to a file</li> </ul> <p>Users who are assigned the Administrator role can also remove alerts, edit alert definitions, and edit alert notification settings.</p>	All functions are available, depending on role.
Tasks	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about tasks</li> <li>• Open task logs</li> <li>• Export task information to a file</li> </ul> <p>If a user assigned the Monitor role has permission to provision storage by using a service class, the user can delete provisioning tasks that the user creates. If the service class does not require administrator approval, the user can also schedule and run the provision tasks that the user creates.</p> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Run tasks</li> <li>• Schedule tasks</li> <li>• Delete tasks</li> </ul>	All functions are available, depending on role.

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Storage systems	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about storage systems</li> <li>• Acknowledge storage system status</li> <li>• Open the management GUI or element manager for a storage system</li> <li>• Export storage system information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Add and remove storage systems</li> <li>• Schedule data collection jobs</li> <li>• Optimize storage tiering</li> <li>• Add storage systems to capacity pools</li> <li>• Administer connections</li> </ul>	Optimize storage tiering
Volumes	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about volumes</li> <li>• Acknowledge volume status</li> <li>• Export volume information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Transform storage volumes</li> <li>• Optimize storage tiering</li> </ul>	Optimize storage tiering
Pools	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about pools</li> <li>• Acknowledge pool status</li> <li>• Export pool information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Add pools to capacity pools</li> <li>• Balance pools</li> <li>• Optimize storage tiering</li> </ul>	<ul style="list-style-type: none"> <li>• Balance pools</li> <li>• Optimize storage tiering</li> </ul>
Shares	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about shares</li> <li>• Export pool information to a file</li> </ul>	All functions are available, depending on role.

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Servers	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about servers</li> <li>• Acknowledge server status</li> <li>• View and collect logs</li> <li>• Export server information to a file</li> </ul> <p>If a user assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to servers.</p> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Add and remove servers</li> <li>• Schedule data collection jobs</li> <li>• Provision storage to servers</li> <li>• Modify Storage Resource agents</li> </ul>	<ul style="list-style-type: none"> <li>• Provision block storage. There are no license restrictions for provisioning file storage.</li> <li>• Optimize storage tiering.</li> </ul>
Hypervisors	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about hypervisors</li> <li>• Acknowledge hypervisor status</li> <li>• Export hypervisor information to a file</li> </ul> <p>If a user who is assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to hypervisors.</p> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Add and remove hypervisors</li> <li>• Schedule data collection jobs</li> <li>• Provision storage to hypervisors</li> <li>• Administer connections</li> </ul>	<p>Provision block storage.</p> <p>There are no license restrictions for provisioning file storage.</p>

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Switches	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about switches</li> <li>• Acknowledge switch status</li> <li>• Open the management GUI or element manager for a switch</li> <li>• Export switch information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Add and remove switches</li> <li>• Schedule data collection jobs</li> <li>• Administer connections</li> <li>• Enable automatic zoning</li> </ul>	Enable automatic zoning
Fabrics	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about fabrics</li> <li>• Acknowledge fabric status</li> <li>• Export fabric information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Add and remove fabrics</li> <li>• Schedule data collection jobs</li> <li>• Administer connections</li> <li>• Enable automatic zoning</li> </ul>	Enable automatic zoning
Service classes	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about service classes</li> <li>• Export service class information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Create service classes</li> <li>• Modify service classes</li> <li>• Delete service classes</li> </ul>	<p>Create block-storage service classes.</p> <p>There are no license restrictions for working with file-storage service classes.</p>

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Capacity pools	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about capacity pools</li> <li>• Export capacity pool information to a file</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Create capacity pools</li> <li>• Modify capacity pools</li> <li>• Delete capacity pools</li> </ul>	All functions are available, depending on role.
Applications and departments	<p>Users who are assigned the Monitor or Administrator role can complete the following monitoring actions:</p> <ul style="list-style-type: none"> <li>• View information about applications and subcomponents</li> <li>• View information about departments and subdepartments</li> <li>• View related resources assigned to applications and departments</li> </ul> <p>Users who are assigned the Administrator role can also complete the following administrative actions:</p> <ul style="list-style-type: none"> <li>• Create applications and departments</li> <li>• Create filters to add resources</li> <li>• Modify filters to change resources</li> <li>• Remove filters</li> <li>• Add resources to applications directly</li> <li>• Remove resources from applications</li> <li>• Add applications as members of other applications</li> <li>• Create departments</li> <li>• Add departments to other departments</li> <li>• Add applications to departments</li> </ul>	All functions are available, depending on role.

## Functions available in the stand-alone GUI based on role and license

When you use the stand-alone GUI, your Tivoli Storage Productivity Center role and product license determine the functions that are available.

Users who are assigned the Administrator role or the Monitor role can use the stand-alone GUI. The functions that are available depend on the role that is assigned to the user:

### Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

### Monitor role

Users who are assigned the Monitor role can view data that is collected by Tivoli Storage Productivity Center, and can create, generate, and save reports.

To use some functions of the stand-alone GUI, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. The following table outlines the functions that are restricted to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. The table mirrors the navigation tree of the stand-alone GUI. The table shows nodes of the navigation tree and indicates whether that function is available with all licenses, or whether the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license is required.

Table 7. Functions with each Tivoli Storage Productivity Center license

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
<b>Administrative Services</b>		
Data Sources <sup>1</sup> <ul style="list-style-type: none"><li>• CIMOM agents</li><li>• Storage Resource agents</li><li>• Storage Subsystems</li><li>• TPC Servers</li><li>• VMWare VI Data Source</li></ul>	<ul style="list-style-type: none"><li>• Yes</li><li>• Yes</li><li>• Yes</li><li>• No</li><li>• Yes</li></ul>	<ul style="list-style-type: none"><li>• Yes</li><li>• Yes</li><li>• Yes</li><li>• Yes</li><li>• Yes</li></ul>
Discovery	Yes	Yes
Configuration	Yes <sup>2</sup>	Yes
<b>IBM Tivoli Storage Productivity Center</b>		
Job Management	Yes	Yes
Reporting <ul style="list-style-type: none"><li>• My Reports (User ID reports)</li><li>• System Reports</li><li>• Rollup Reports</li><li>• Data Source Reports</li></ul>	Yes See Table 8 on page 27 for a list of available reports.	Yes See Table 8 on page 27 for a list of available reports.
Topology Viewer	Yes	Yes
Monitoring (data collection) <ul style="list-style-type: none"><li>• Probe Alert Configurations</li><li>• TPC Server Probes</li></ul>	<ul style="list-style-type: none"><li>• Yes</li><li>• No</li></ul>	<ul style="list-style-type: none"><li>• Yes</li><li>• Yes</li></ul>
Storage Resource Group Management	Yes <sup>5</sup>	Yes
Analytics <ul style="list-style-type: none"><li>• Configuration History</li><li>• Configuration Analysis</li></ul>	<ul style="list-style-type: none"><li>• No</li><li>• No</li></ul>	<ul style="list-style-type: none"><li>• Yes</li><li>• Yes</li></ul>
Alerting:		

Table 7. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
<ul style="list-style-type: none"> <li>• Authentication Configuration Alerts</li> <li>• Alert Log</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes<sup>3</sup></li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> </ul>
<b>Data Manager</b>		
Monitoring: (data collection)		
<ul style="list-style-type: none"> <li>• Groups</li> <li>• Pings</li> <li>• Scans</li> <li>• Profiles</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• No</li> <li>• No</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>
Alerting	Yes <sup>4</sup>	Yes
Policy Management	No	Yes
Reporting	Yes  Not all reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. See Table 8 on page 27 for a list of available reports.	Yes
<b>Data Manager for Databases</b>	No	Yes
<b>Data Manager for Chargeback</b>	No	Yes
<b>Disk Manager</b>		
Storage Subsystem	Yes	Yes
Monitoring (data collection)		
<ul style="list-style-type: none"> <li>• Groups</li> <li>• Jobs</li> <li>• Subsystem Performance Alert Configurations</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>
Alerting	Yes	Yes
Reporting		
<ul style="list-style-type: none"> <li>• Groups</li> <li>• Storage Subsystems</li> <li>• Storage Subsystem Performance</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>
<b>Fabric Manager</b>		
Monitoring (data collection)		
<ul style="list-style-type: none"> <li>• Groups</li> <li>• Switch Performance Alert Configurations</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> </ul>
Fabrics	Yes	Yes



Table 7. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Alerting	Yes	Yes
• Switch Performance	Yes	Yes
<b>Replication Manager</b>	Yes	Yes
<p><b>Note:</b></p> <ol style="list-style-type: none"> <li>You cannot associate TPC Servers with Tivoli Storage Productivity Center in the Tivoli Storage Productivity Center or Storage Productivity Center Select license. VMware Data Sources are associated with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. TPC Servers are associated with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.</li> <li>The following nodes under <b>Configuration</b> are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> <li>• License Keys</li> <li>• Quota and Constraint e-mail Address Rules</li> <li>• Resource History Retention for Databases</li> <li>• Removed Resource Retention for Databases</li> <li>• Configuration History Settings</li> </ul> </li> <li>The following nodes under <b>Alert Log</b> are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> <li>• Directory</li> <li>• User</li> <li>• OS User Group</li> <li>• Configuration Analysis</li> </ul> </li> <li>The following node under <b>Data Manager &gt; Alerting</b> is not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: Directory Alerts. The node for Other NAS Alerts is available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license.</li> <li>You cannot create a storage resource group for a type of capacity pool.</li> </ol>		

Table 8. Available reports with each IBM Tivoli Storage Productivity Center license

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Batch reports	Yes <sup>1, 2</sup>	Yes <sup>1</sup>
System Reports:		
• Data	• Yes <sup>4</sup>	• Yes
• Fabric	• Yes	• Yes
• Disk	• Yes	• Yes
Rollup Reports	No	Yes
Data Source Reports	Yes	Yes
Asset reports		

Table 8. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
<ul style="list-style-type: none"> <li>• By Cluster</li> <li>• By Computer</li> <li>• By Hypervisor</li> <li>• By IBM Storwize V7000 Unified/IBM SONAS</li> <li>• By OS Type</li> <li>• By Storage Subsystem</li> <li>• System-wide</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes<sup>3</sup></li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>
Availability reports	Yes	Yes
TPC-wide Storage Space reports:		
<ul style="list-style-type: none"> <li>• Disk Space</li> <li>• File System Space</li> <li>• Consumed File System Space</li> <li>• Available File System Space</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>
Usage reports	No	Yes (these reports are a result of a scan)
Usage Violation reports	No	Yes (these reports are a result of a scan)
Backup reports:	No	Yes (these reports are a result of a scan)
Monitored Computer Storage Space reports	Yes	Yes
Storage Subsystem reports	Yes	Yes
Storage Subsystem Performance reports	Yes	Yes
Switch Performance reports	Yes	Yes

Table 8. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
<p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. A Storage Resource agent must be deployed on the computer where you want to run batch reports.</li> <li>2. The following batch reports are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> <li>• Rollup</li> <li>• Usage</li> <li>• Usage Violations</li> <li>• Backup</li> </ul> </li> <li>3. The following <b>Asset &gt; System-wide</b> reports are available with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license: <ul style="list-style-type: none"> <li>• Monitored Directories</li> <li>• Unmanaged Computers</li> <li>• Users</li> <li>• OS User Groups</li> </ul> </li> <li>4. The following Data system reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> <li>• Disk Space Summary</li> <li>• Disk Defects</li> <li>• Computer Storage Availability</li> <li>• Computer Disk Space</li> <li>• Available File System Space</li> </ul> </li> </ol>		



---

## Chapter 3. Product overview

IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance, and resources. It can reduce the complexity of managing a storage environment by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, replication services, and capacity management.

Because it is policy-based, Tivoli Storage Productivity Center can detect potential problems and automatically adjust storage and resources that are based on the policies and actions that you define. For example, it can notify you when a server or storage system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues that are related to your stored data, you can use Tivoli Storage Productivity Center to prevent unnecessary system and application downtime.

Tivoli Storage Productivity Center provides the following benefits and functions:

- Simplifies the management of storage infrastructures
- Manages, administers, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached devices
- Monitors, manages, and controls (through zones) SAN fabric components
- Manages the capacity utilization and availability of file systems and databases

---

## Architecture

The IBM Tivoli Storage Productivity Center consists of several components that form the infrastructure of its storage-management functions.

### Data server

This component is the control point for product scheduling functions, configuration, event information, reporting, and graphical user interface (GUI) support. It coordinates communication with and data collection from agents that scan file systems and databases to gather storage demographics and populate the database with results. Automated actions can be defined to perform file system extension, data deletion, and Tivoli Storage Productivity Center backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for GUI user interface functions. It also includes functions that schedule data collection and discovery for the Device server.

### Device server

This component discovers, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

### Replication server

This component coordinates communication and processes tasks that are related to replication and Tivoli Storage Productivity Center for Replication. It also provides the infrastructure for hosting the Tivoli Storage Productivity Center for Replication GUI.

**Web server**

Tivoli Storage Productivity Center uses the WebSphere Application Server as the web application server to host its web-based GUI, storage management API for cloud, and VMware vSphere web client extension.

**Jazz for Service Management and Tivoli Common Reporting (optional)**

These components enable Tivoli Storage Productivity Center reports in the web-based GUI. Tivoli Common Reporting includes IBM Cognos reporting software. You can use this software to create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors.

For large storage environments where one server is not sufficient to manage all the Tivoli Storage Productivity Center components, you can install the product on multiple servers. In this type of environment, Jazz for Service Management and Tivoli Common Reporting can run on a separate server.

**Database**

A single database instance serves as the repository for all Tivoli Storage Productivity Center components.

**Agents**

Storage Resource agents, CIM agents, and Out of Band fabric agents gather host, application, storage system, and SAN fabric information and send that information to the Data Server or Device server.

**GUI** Tivoli Storage Productivity Center provides two graphical user interfaces for managing the storage infrastructure in an enterprise environment: a stand-alone GUI and the web-based GUI. Each GUI provides different functions for working with monitored resources. To view a comparison of the functions that are available in each GUI, see “Available functions in the interfaces” on page 34.

**CLI** Use the command-line interface (CLI) to issue commands for key Tivoli Storage Productivity Center functions.

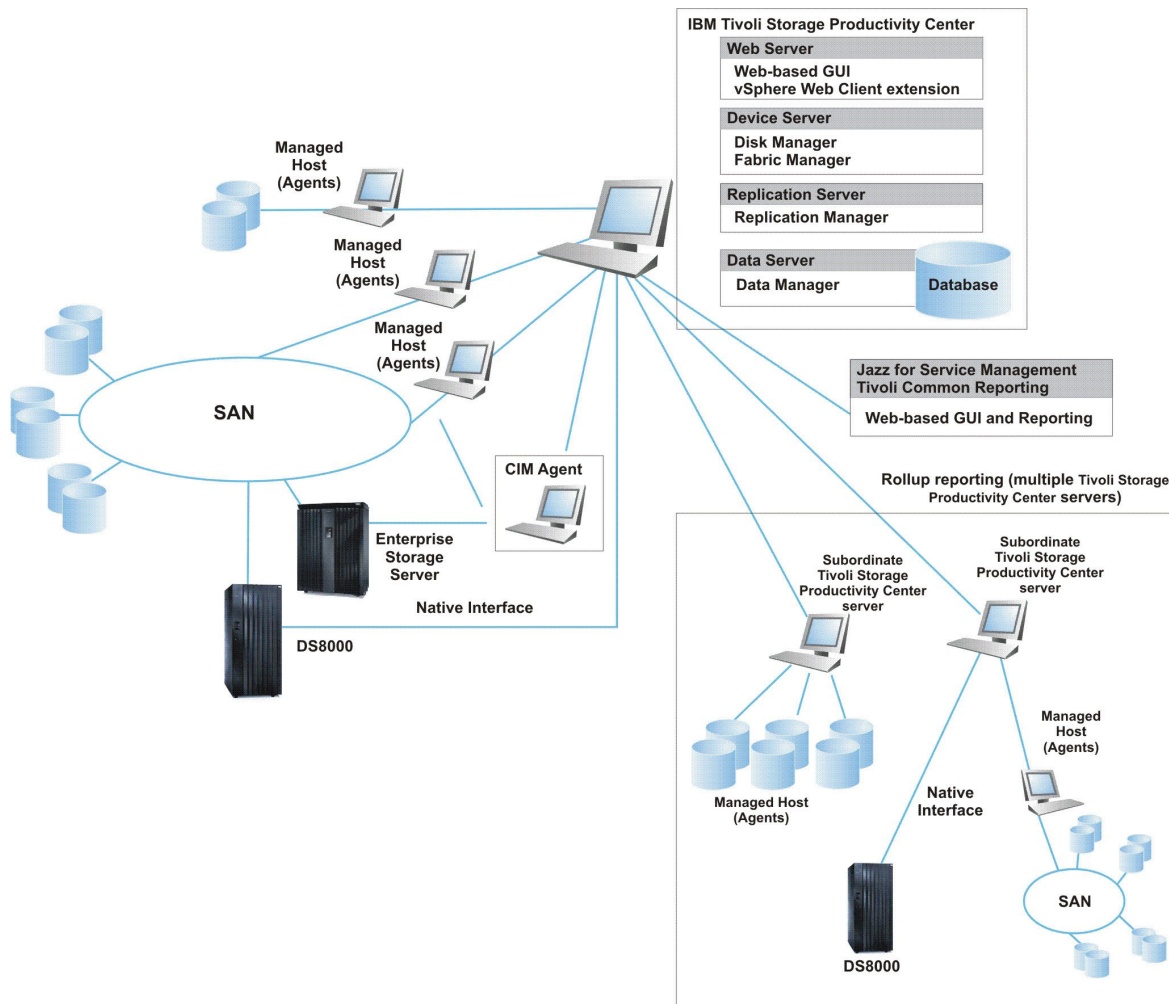


Figure 1. Tivoli Storage Productivity Center components

## Interfaces for Tivoli Storage Productivity Center

Tivoli Storage Productivity Center provides multiple user interfaces for managing the storage infrastructure in an enterprise environment.

### Web-based GUI

This interface runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can also report on the condition, capacity, and performance resources in the Tivoli Common Reporting interface. In 5.2 and later, the web-based GUI includes many of the functions for managing storage, disks, and fabrics. You can access this interface from anywhere that you have a web browser and connectivity to a network. For more information about the functions that are available in this GUI, see “Available functions in the interfaces” on page 34.

### IBM Cognos reporting interface

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in the Cognos reporting software. The

Cognos software is included in Tivoli Common Reporting. For more information about reporting in the web-based GUI, see Chapter 7, “Reporting,” on page 483.

### Stand-alone GUI

The stand-alone GUI is the same, Java based interface that was available in Tivoli Storage Productivity Center v4.2.2.x and earlier. In version 5.2 and later, many of the functions for monitoring and managing storage in this interface were moved to the web-based GUI. For more information about the functions that were moved from this GUI to the web-based GUI, see “Functions that were moved to the web-based GUI in V5.2 and later” on page 39.

### Command-line interface

Use this interface to run Tivoli Storage Productivity Center commands from a command prompt. For more information about CLI commands, see the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference*.

### Tivoli Storage Productivity Center for Replication interface

Tivoli Storage Productivity Center for Replication provides a GUI that runs in a web browser and a command-line interface. Use these interfaces to manage data replication and disaster recovery. For more information about these interfaces, see “Interfaces for Tivoli Storage Productivity Center for Replication” on page 49. For more information about the features that are available in Tivoli Storage Productivity Center for Replication, see Chapter 6, “Managing replication,” on page 303.

### Related concepts:

Chapter 5, “Managing resources,” on page 127

Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

## Available functions in the interfaces

Tivoli Storage Productivity Center provides the web-based GUI and stand-alone GUI for managing storage resources. In V5.2, V5.2.1, and V5.2.2, many of the functions for managing storage were moved to the web-based GUI. Determine which GUI to use based on the tasks that you must complete.

## Configuring and administering

Use the following list to determine the configuring and administering tasks that you can complete in each GUI.

Table 9. Available configuring and administering tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Discovering resources for monitoring	Yes <sup>1</sup>	Yes <sup>1</sup>
Adding resources for monitoring	Yes <sup>2</sup>	Yes <sup>2</sup>
Removing resources (data sources)	Yes <sup>3</sup>	Yes <sup>3</sup>



Table 9. Available configuring and administering tasks in each GUI (continued)

Tasks	Available in web-based GUI	Available in stand-alone GUI
Administering resources (data sources)	Yes: <ul style="list-style-type: none"> <li>Viewing information about data sources</li> <li>Viewing information about the resources that are managed by data sources</li> <li>Viewing and modifying connections</li> <li>Testing connections</li> <li>Updating credentials</li> </ul>	Yes, but only for the following actions and data sources: <ul style="list-style-type: none"> <li>Viewing information about data sources</li> <li>Viewing information about the resources that are managed by data sources</li> <li>Testing connections: storage systems, TPC servers</li> <li>Launching GUI: TPC servers</li> </ul>
Administering Tivoli Storage Productivity Center services	Yes	No
Administering servers and Storage Resource agents	Yes	Yes, but only to view logs
Configuring Tivoli Storage Productivity Center	Yes: <ul style="list-style-type: none"> <li>Assigning roles to user groups</li> <li>Modifying the authentication repository</li> <li>Specifying how long to retain the data that is collected about resources</li> <li>Specifying how long to retain log files</li> <li>Specifying how to receive alerts</li> <li>Upgrading Storage Resource agents</li> <li>Setting trace levels for component servers</li> </ul>	All other configuration tasks
Viewing and managing tasks	Yes, for the following types of tasks: <ul style="list-style-type: none"> <li>Provisioning</li> <li>Tiering</li> <li>Transforming</li> <li>Balancing pools</li> </ul>	All other tasks
Viewing and managing data collection jobs	Yes, for the following types data collection jobs: <ul style="list-style-type: none"> <li>Probes</li> <li>Performance monitors</li> <li>Agent upgrades</li> </ul>	All other jobs

Table 9. Available configuring and administering tasks in each GUI (continued)

Tasks	Available in web-based GUI	Available in stand-alone GUI
Opening the management GUIs for resources	You can open the start page of the management GUI for a resource. The GUI must be web-based or use Java Web Start.	No
<p><b>Notes for configuring and administering tasks:</b></p> <ol style="list-style-type: none"> <li>In the web-based GUI, the discovery process is completed automatically when you add resources. In the stand-alone GUI, you must run a discovery job to find resources before you can add them.</li> <li>In the web-based GUI, you can add the following resources for monitoring: storage systems, servers, hypervisors, fabrics, and switches. In the stand-alone GUI, you can add Tivoli Storage Productivity Center servers as subordinate servers.</li> <li>In the web-based GUI, you can remove the following resources: servers, hypervisors, storage systems, switches, fabrics, NAS filers, and CIM agents for storage systems. In the stand-alone GUI, you can remove the following resources: CIM agents, VMware vSphere servers, out of band fabric agents, and Tivoli Storage Productivity Center servers as subordinate servers.</li> </ol>		

## Monitoring

Use the following list to determine the monitoring tasks that you can complete in each GUI.

Table 10. Available monitoring tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Scheduling data collection	Yes: <ul style="list-style-type: none"> <li>Probes</li> <li>Performance monitors</li> </ul>	Yes: <ul style="list-style-type: none"> <li>Pings</li> <li>Scans</li> </ul>
Creating groups of resources	No	Yes
Creating storage resource groups	No	Yes
Viewing information about resources	Yes, except information about databases	Yes
Viewing the following information about resources: <ul style="list-style-type: none"> <li>Status</li> <li>Capacity (including used and available space)</li> <li>Relationships between resources</li> <li>Data paths</li> </ul>	<ul style="list-style-type: none"> <li>Yes</li> <li>Yes</li> <li>Yes</li> <li>Yes</li> </ul>	<ul style="list-style-type: none"> <li>Yes</li> <li>Yes</li> <li>Yes</li> <li>Yes</li> </ul>
Exporting information about resources to a file	Yes	Yes
Viewing a visual topology of resources	Yes, for storage systems, servers, and hypervisors	Yes
Viewing a visual data path of resources	Yes	Yes

Table 10. Available monitoring tasks in each GUI (continued)

Tasks	Available in web-based GUI	Available in stand-alone GUI
Creating applications and departments	Yes	No
Viewing information about applications and departments	Yes	No
Viewing information about the resources that are related to applications and departments	Yes	No
Exporting information about applications and departments to a file	Yes	No
Troubleshooting performance issues	Yes	Yes, but only for viewing the performance reports that were available in previous versions.
Running analytics and configuring storage: <ul style="list-style-type: none"> <li>• Configuration history</li> <li>• Configuration analysis</li> <li>• Optimizing and balancing storage</li> <li>• Provisioning storage</li> <li>• Cloud configuration</li> </ul>	<ul style="list-style-type: none"> <li>• No</li> <li>• No</li> <li>• Yes</li> <li>• Yes</li> <li>• Yes</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes</li> <li>• No</li> <li>• No</li> <li>• No</li> </ul>
Managing volumes for storage systems <sup>1</sup> : <ul style="list-style-type: none"> <li>• Viewing information about volumes</li> <li>• Creating volumes</li> <li>• Deleting volumes</li> <li>• Assigning volumes to host ports</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes, when provisioning storage</li> <li>• No</li> <li>• Yes, when provisioning storage</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• No</li> <li>• No</li> <li>• No</li> </ul>
Managing virtual disks for storage systems <sup>1</sup> : <ul style="list-style-type: none"> <li>• Viewing information about virtual disks</li> <li>• Creating virtual disks</li> <li>• Deleting virtual disks</li> <li>• Assigning volumes to host ports</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes, when provisioning storage</li> <li>• No</li> <li>• Yes, when provisioning storage</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• No</li> <li>• No</li> <li>• No</li> </ul>
Managing fabrics	Yes, but only for setting a zoning policy to enable automatic zoning. When automatic zoning is enabled, Tivoli Storage Productivity Center can create zones during storage provisioning.	Yes, but only for viewing zone configuration information
Troubleshooting performance issues	Yes	Yes, by using reports
Managing zones and zone sets	No	No
Creating storage policies	No	Yes
Charging back for storage usage	No	Yes

Table 10. Available monitoring tasks in each GUI (continued)

Tasks	Available in web-based GUI	Available in stand-alone GUI
Profile management	Yes, managing service classes	No
<b>Restriction:</b> <sup>1</sup> Certain actions for provisioning storage and managing virtual disks are no longer available in Tivoli Storage Productivity Center. To complete actions that were not added to the web-based GUI, use the element manager or management GUI for the related resource.		

## Alerting

Use the following list to determine the alerting tasks that you can complete in each GUI.

Table 11. Available alerting tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Creating alerts	Yes	Yes
Viewing alert conditions that are detected on monitored resources	Yes	Yes
Viewing replication alerts	Yes	Yes
Acknowledging the status of alerts	Yes	Yes
Removing alerts	Yes	Yes

## Reporting

Use the following list to determine the reporting tasks that you can complete and information that you can view in each GUI.

### Reporting tasks

Table 12. Available reporting tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Accessing the Cognos reports GUI	Yes	No
Saving reports	Yes	Yes
Scheduling reports	Yes	Yes

### Reporting information

Table 13. Available reporting information in each GUI

Information	Available in web-based GUI	Available in stand-alone GUI
Resource properties	Yes	Yes
Performance	Yes	Yes
Capacity (including used and available space)	Yes	Yes
Historical capacity	Yes	Yes
Storage tiering	Yes	No

Table 13. Available reporting information in each GUI (continued)

Information	Available in web-based GUI	Available in stand-alone GUI
Policy violations	No	Yes
Rollup servers	No	Yes
File statistics	No	Yes
Data sources	No	Yes
Databases	No	Yes

## Functions that were moved to the web-based GUI in V5.2 and later

In Tivoli Storage Productivity Center V5.2 and later, a number of product functions were moved from the stand-alone GUI to the web-based GUI. Use the following list to determine which function were moved.

Table 14. Functions that were moved from the stand-alone GUI to the web-based GUI

Tasks	Moved to the web-based GUI
Configuring <sup>1</sup>	<ul style="list-style-type: none"> <li>• Assigning roles to user groups</li> <li>• Specifying how long to retain the data that is collected about resources</li> <li>• Specifying how long to retain log files</li> <li>• Specifying how to receive alerts</li> <li>• Upgrading Storage Resource agents</li> <li>• Setting trace levels for component servers</li> </ul>
Discovering resources	<ul style="list-style-type: none"> <li>• Hypervisors (VMware virtual infrastructure)</li> <li>• Storage systems</li> <li>• Switches</li> <li>• Fabrics</li> </ul> <p>In the web-based GUI, discovering resources is done automatically when you add resources in the Add resource wizard.</p>
Adding resources (data sources) for monitoring	<ul style="list-style-type: none"> <li>• Servers</li> <li>• Hypervisors (VMware virtual infrastructure)</li> <li>• Storage systems</li> <li>• Switches</li> <li>• Fabrics</li> <li>• CIM agents for storage systems, switches, fabrics, and NAS filers</li> </ul>

Table 14. Functions that were moved from the stand-alone GUI to the web-based GUI (continued)

Tasks	Moved to the web-based GUI
Removing resources (data sources)	<ul style="list-style-type: none"> <li>• Servers</li> <li>• Hypervisors</li> <li>• Storage systems</li> <li>• Switches</li> <li>• Fabrics</li> <li>• NAS filers</li> </ul>
Administering resources (data sources)	<ul style="list-style-type: none"> <li>• Viewing and modifying connections</li> <li>• Testing connections</li> <li>• Updating credentials</li> </ul>
Administering servers and Storage Resource agents	<ul style="list-style-type: none"> <li>• Deploying agents</li> <li>• Disabling agents</li> <li>• Enabling agents</li> <li>• Upgrading agents</li> <li>• Updating user credentials</li> <li>• Viewing logs</li> <li>• Collecting agent logs</li> <li>• Testing connections</li> <li>• Enabling and disabling fabric functions</li> <li>• Enabling and disabling running scripts on agent</li> <li>• Removing servers</li> </ul>
Starting element managers (management GUIs) in context <sup>2</sup>	You can open the start page of an element that is web-based or uses Java Web Start.
Starting Tivoli Storage Productivity Center for Replication in context <sup>3</sup>	Yes, from the Alerts page. If a replication alert includes a link to a storage system, you can click the link to start Tivoli Storage Productivity Center for Replication.
Scheduling data collection	Scheduling the following types of jobs: <ul style="list-style-type: none"> <li>• Probes</li> <li>• Performance monitors</li> </ul>
Viewing and managing jobs and tasks	Viewing and managing the following types of jobs and tasks: <ul style="list-style-type: none"> <li>• Probes</li> <li>• Performance monitors</li> <li>• Agent upgrades</li> <li>• Provisioning (task)</li> <li>• Optimizing (task)</li> </ul>

Table 14. Functions that were moved from the stand-alone GUI to the web-based GUI (continued)

Tasks	Moved to the web-based GUI
Defining alerts for configuration changes and violations of performance thresholds	Yes, except for some configuration change conditions related to computers, hypervisors, directories, fabrics, and endpoint devices. All other conditions are available in the web-based GUI.
Troubleshooting performance	Yes
Viewing a visual topology of resources	Yes, for viewing data paths
Analytics <ul style="list-style-type: none"> <li>• Provisioning storage</li> <li>• Optimizing storage</li> <li>• Cloud Configuration               <ul style="list-style-type: none"> <li>– Assigning storage to tiers</li> <li>– Creating and modifying service classes</li> <li>– Creating and modifying capacity pools</li> </ul> </li> </ul>	Yes
Viewing information about managed disk groups (pools)	Yes
Managing volumes and virtual disks for storage systems <sup>4</sup> : <ul style="list-style-type: none"> <li>• Viewing information about volumes and virtual disks</li> <li>• Creating volumes and virtual disks</li> <li>• Assigning volumes to host ports</li> </ul>	<ul style="list-style-type: none"> <li>• Yes</li> <li>• Yes, when provisioning storage</li> <li>• Yes, when provisioning storage</li> </ul>
Managing fabrics	Setting a zoning policy to enable automatic zoning. When automatic zoning is enabled, Tivoli Storage Productivity Center can create zones during storage provisioning.
<b>Restrictions:</b> <ul style="list-style-type: none"> <li>• <sup>1</sup> The ability to migrate Data agents and Fabric agents is no longer available in either GUI. To monitor a resource where a Data agent or Fabric agent is installed, you must uninstall that agent and manually deploy a Storage Resource agent.</li> <li>• <sup>2</sup> If the firmware level for a DS8000 is earlier than R6.2, you must complete extra configuration steps before you can open its management GUI. For information, go to “Opening the management GUI for a resource” on page 284.</li> <li>• <sup>3</sup> The ability to open pages in the Tivoli Storage Productivity Center for Replication GUI from the <b>Replication Management</b> node in the stand-alone GUI is no longer available. In 5.2 and later, you must start the Tivoli Storage Productivity Center for Replication GUI as a stand-alone application in a web browser. The exception is starting the GUI from the Alerts page as described in the table.</li> <li>• <sup>4</sup> Certain actions for managing volumes and virtual disks are no longer available in Tivoli Storage Productivity Center (for example, deleting). To complete actions that were not added to the web-based GUI, use the storage management GUI for the related resource.</li> </ul>	

---

## Starting Tivoli Storage Productivity Center

Tivoli Storage Productivity Center provides two graphical user interface (GUI) applications for managing and monitoring the resources in a storage environment: a web-based GUI and a stand-alone GUI. You can start these GUIs on the Tivoli Storage Productivity Center server or on a remote system.

Each GUI provides different functions for working with monitored resources.

### Starting the Tivoli Storage Productivity Center stand-alone GUI

The stand-alone GUI contains functions for monitoring the condition of storage resources, and all the tools for managing data, disk, fabric, and replication.

How you start the stand-alone GUI depends on whether the stand-alone GUI component is installed on your system.

- On the Tivoli Storage Productivity Center server, or on a remote system where the stand-alone GUI component is installed, you start the stand-alone GUI by running a batch file or a shell script. For Windows operating systems, you run the batch file `tpc.bat`, which you can run from the **Start** menu. For AIX or Linux operating systems, you run the shell script `TPCD.sh`.
- If the stand-alone GUI component is not installed on a remote system, you use a Java Web Start application to download and start the stand-alone GUI.

#### Starting the Tivoli Storage Productivity Center stand-alone GUI by using Java Web Start

Start the stand-alone Tivoli Storage Productivity Center GUI by using Java Web Start to interact remotely with Tivoli Storage Productivity Center running on the server.

Before you start Tivoli Storage Productivity Center, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>. In the **Agents, Servers and Browser** column, click the version of Tivoli Storage Productivity Center that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

On remote systems where the stand-alone GUI component is not installed, you can start the stand-alone Tivoli Storage Productivity Center GUI using Java Web Start. To start the stand-alone GUI by using Java Web Start, you click a web link that downloads a Java Network Launching Protocol (JNLP) file from the server. The JNLP file specifies all the files needed by the application. The Java Web Start Launcher on the remote system opens the JNLP file, downloads and caches all the required files, and starts the stand-alone GUI in a new window.

1. Start a web browser, and type the following address for the Java Web Start page.

```
http://hostname:port/ITSRM/app/welcome.html
```

In the preceding address, specify the following values:

*hostname*

The Tivoli Storage Productivity Center server. You can specify the *hostname* as an IP address or a Domain Name System (DNS) name. To verify your host name, ask your Tivoli Storage Productivity Center administrator.



*port*


The port number for the Device server. The default port number for connecting to the Device server using the HTTP protocol is 9550. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your Tivoli Storage Productivity Center administrator for the correct port number.

2. The Tivoli Storage Productivity Center GUI requires an IBM Java Runtime Environment 1.6.0. Links are provided on the Java Web start page for you to download the IBM Java Runtime Environment 1.6.0 for Windows, Linux, or AIX. If an IBM Java Runtime Environment 1.6.0 is not already installed on the system, click the link to download the IBM Java Runtime Environment 1.6.0 for your operating system. On Windows, download a self-extracting executable. On Linux, download an RPM file. On AIX, download an executable JAR file. Install the JRE from the file that you downloaded before you start the stand-alone GUI. If you are prompted to install the JRE as the system JRE or to overwrite the current system JRE, click **No**.
3. From the Java Web Start page, click **TPC GUI (Java Web Start)**  
The JNLP file is downloaded.

**Important:** Depending on your browser and operating system and their default settings for your browser and operating system) you might need to:

- Verify to your browser that the `tpcgui.jnlp` file is safe to download
- Indicate that JNLP files are to be opened by the Java Web Start Launcher

For more information, go to the product documentation at

 [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_tbs\\_gui\\_does\\_not\\_open.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_tbs_gui_does_not_open.html).

4. The stand-alone Tivoli Storage Productivity Center GUI opens in a separate window. To log on to Tivoli Storage Productivity Center, enter the following information and click **OK**.
  - a. Enter the user ID and password that was defined for the installation of Tivoli Storage Productivity Center.
  - b. If the server field does not contain the address and port number of for the server, specify the following values for the *hostname* and *port*.

*hostname*

The Tivoli Storage Productivity Center server. You can specify the *hostname* as an IP address or a Domain Name System (DNS) name. To verify your host name, ask your Tivoli Storage Productivity Center administrator.

*port*

The port number for the Data server. The default port number for connecting to the Data server is 9549. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation.

## Starting the Tivoli Storage Productivity Center stand-alone GUI on a system where the GUI component is installed

On the Tivoli Storage Productivity Center server, or on a remote system where the stand-alone GUI component is installed, you start the Tivoli Storage Productivity Center stand-alone GUI using the batch file `tpc.bat` on a Windows operating system or the shell script `TPCD.sh` on an AIX and Linux operating system.

A typical installation of Tivoli Storage Productivity Center on the server includes the stand-alone GUI component. A remote system can also have the stand-alone GUI component installed. On a system where the stand-alone GUI component is installed, use the following instructions for starting the stand-alone GUI. If you are on a remote system where the stand-alone GUI component is not installed, start the stand-alone GUI by using Java Web Start instead.

1. Start the Tivoli Storage Productivity Center GUI on Windows operating systems.
  - Open the Tivoli Storage Productivity Center stand-alone GUI. . If you are on a remote system, the **Start** menu or Start page for the stand-alone GUI might not be available. In this case, run the `tpc.bat` file directly. The default location for the batch file is `C:\Program Files\IBM\TPC\gui\tpc.bat`.
  - On AIX or Linux, type the following path and command at the command line.  
`/opt/IBM/TPC/gui/TPCD.sh`
2. From the Tivoli Storage Productivity Center GUI logon window, log on to Tivoli Storage Productivity Center.
  - a. Enter your user ID and password.
  - b. If the server field does not contain the address of the server, enter the information in the format `hostname:port`. Specify the following values for the `hostname` and `port`:  
  
*hostname*  
The Tivoli Storage Productivity Center server. You can specify the *hostname* as an IP address or a Domain Name System (DNS) name.  
  
*port*  
The port number for the Data server. The default port number for connecting to the Data server is 9549. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation.
  - c. Click **OK**.

## Starting the Tivoli Storage Productivity Center web-based GUI

You can start the web-based Tivoli Storage Productivity Center GUI by opening a web browser and entering a web address for the Tivoli Storage Productivity Center logon page. For example, you might enter `http://storage.example.com:9568/srm`.

Before you start IBM Tivoli Storage Productivity Center, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>. In the **Agents, Servers and Browser** column, click the version of Tivoli Storage Productivity Center that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

The web-based GUI provides quick access to pages that you can use to monitor the condition, capacity, and relationships of the resources within your storage environment. Start the web-based GUI if you are interested in monitoring your storage environment, but do not need the full set of administrative tools available in the stand-alone GUI.

1. On a server running the Windows operating system, start **TPC Web-based GUI**. If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

`http://hostname:port/srm`

In the preceding address, specify the following values:

*hostname*

The Tivoli Storage Productivity Center server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

*port*

The port number for Tivoli Storage Productivity Center. The default port number for connecting to Tivoli Storage Productivity Center by using the HTTP protocol is 9568. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your Tivoli Storage Productivity Center administrator for the correct port number.

**Tip:** If you have a non-default port, check the value of the `WC_defaulthost` property in `TPC_installation_directory/web/conf/portdef.props` file.

You might be redirected from the address that you enter to another address and port that provides secure access using the HTTPS protocol. This page is the Tivoli Storage Productivity Center logon page.

2. From the Tivoli Storage Productivity Center logon page, type your user name and password and click **Log in**. The web-based Tivoli Storage Productivity Center GUI opens in the browser.

## Starting Tivoli Storage Productivity Center for Replication

You can start the IBM Tivoli Storage Productivity Center for Replication GUI by opening a web browser and entering a web address for the Tivoli Storage Productivity Center for Replication logon page.

Before you start Tivoli Storage Productivity Center for Replication, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>. In the **Agents, Servers and Browser** column, click the version of Tivoli Storage Productivity Center that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

The Tivoli Storage Productivity Center for Replication GUI provides a single point of control to configure, manage, and monitor copy services. Whether you start it on the server or on a remote system, the GUI is web-based and is displayed in a browser.

1. On a server that is running on the Windows operating system, click **Start > Programs > IBM Tivoli Storage Productivity Center > TPC Replication Manager GUI** to open a browser to the correct address. If you are not using a server that is running on the Windows operating system, start a web browser and enter the following address in the address field. The address is case-sensitive.

`https://hostname:port/CSM`

In the preceding address, specify the following values:

### *hostname*

The Tivoli Storage Productivity Center for Replication server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

### *port*

The port number for Tivoli Storage Productivity Center for Replication. The default port number for connecting to Tivoli Storage Productivity Center for Replication using the HTTPS protocol is 9559. However, this port number might be different for your site. For example:

- If you upgraded from Tivoli Storage Productivity Center for Replication version 3.3 or earlier, the default HTTPS port is 9443.
- If you upgraded from Tivoli Storage Productivity Center for Replication version 3.4 or later, the default HTTPS port is 3443.
- If you are using IBM WebSphere Application Server for z/OS®, the default HTTPS port is 9443.
- If you are using IBM WebSphere Application Server OEM Edition for z/OS, the default HTTPS port is 32209.

You can find the port number for the Replication server in the `install_root/wlp/usr/servers/replicationServer/properties/portdef.props` file. The default port for the HTTP protocol is defined by the `WC_defaulthost_secure` property in the file.

You can find the port number for the web-based GUI in the `install_root/eWAS/profiles/webServerProfile/properties/portdef.props` file. The default port for the HTTP protocol is defined by the `WC_defaulthost` property in the file.

2. From the Tivoli Storage Productivity Center for Replication logon page, type your user name and password and click **Login**. The web-based Tivoli Storage Productivity Center for Replication GUI opens in the browser.

---

## Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center includes Tivoli Storage Productivity Center for Replication. Tivoli Storage Productivity Center for Replication manages copy services in storage environments. *Copy services* are features that are used by storage systems such as System Storage DS8000 to configure, manage, and monitor data-copy functions. Copy services include IBM FlashCopy, Metro Mirror, Global Mirror, and Metro Global Mirror.

Tivoli Storage Productivity Center for Replication provides copy services management for the storage systems and session types that are listed in the following table:

Storage system	Session type
System Storage DS8000	<ul style="list-style-type: none"> <li>• Basic HyperSwap</li> <li>• FlashCopy</li> <li>• Metro Mirror Single Direction</li> <li>• Metro Mirror Failover/Failback</li> <li>• Metro Mirror Failover/Failback with Practice</li> <li>• Global Mirror Single Direction</li> <li>• Global Mirror Either Direction with Two-Site Practice</li> <li>• Global Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback with Practice</li> <li>• Metro Global Mirror</li> <li>• Metro Global Mirror with Practice</li> <li>• Metro Mirror - Metro Mirror (This session type is available only for storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.)</li> </ul>
System Storage DS6000	<ul style="list-style-type: none"> <li>• Basic HyperSwap</li> <li>• FlashCopy</li> <li>• Metro Mirror Single Direction</li> <li>• Metro Mirror Failover/Failback</li> <li>• Metro Mirror Failover/Failback with Practice</li> <li>• Global Mirror Single Direction</li> <li>• Global Mirror Either Direction with Two-Site Practice</li> <li>• Global Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback with Practice</li> </ul>
TotalStorage Enterprise Storage Server® Model 800	<ul style="list-style-type: none"> <li>• Basic HyperSwap</li> <li>• FlashCopy</li> <li>• Metro Mirror Single Direction</li> <li>• Metro Mirror Failover/Failback</li> <li>• Metro Mirror Failover/Failback with Practice</li> <li>• Global Mirror Single Direction</li> <li>• Global Mirror Either Direction with Two-Site Practice</li> <li>• Global Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback with Practice</li> <li>• Metro Global Mirror (only H1 site)</li> <li>• Metro Global Mirror with Practice (only H1 site)</li> </ul>

Storage system	Session type
SAN Volume Controller	<ul style="list-style-type: none"> <li>• FlashCopy</li> <li>• Metro Mirror Single Direction</li> <li>• Metro Mirror Failover/Failback</li> <li>• Metro Mirror Failover/Failback with Practice</li> <li>• Global Mirror Single Direction</li> <li>• Global Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback with Practice</li> <li>• Global Mirror Failover/Failback with Change Volumes</li> </ul>
Storwize V3500	<ul style="list-style-type: none"> <li>• FlashCopy</li> </ul>
Storwize V3700 and Storwize V7000	<ul style="list-style-type: none"> <li>• FlashCopy</li> <li>• Metro Mirror Single Direction</li> <li>• Metro Mirror Failover/Failback</li> <li>• Metro Mirror Failover/Failback with Practice</li> <li>• Global Mirror Single Direction</li> <li>• Global Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback with Practice</li> <li>• Global Mirror Failover/Failback with Change Volumes</li> </ul>
Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)	<ul style="list-style-type: none"> <li>• FlashCopy</li> <li>• Metro Mirror Single Direction</li> <li>• Metro Mirror Failover/Failback</li> <li>• Metro Mirror Failover/Failback with Practice</li> <li>• Global Mirror Single Direction</li> <li>• Global Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback with Practice</li> <li>• Global Mirror Failover/Failback with Change Volumes</li> </ul>
The XIV system	<ul style="list-style-type: none"> <li>• Snapshot</li> <li>• Metro Mirror Failover/Failback</li> <li>• Global Mirror Failover/Failback</li> </ul>

## Key components of Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication consists of several key components.

### Tivoli Storage Productivity Center for Replication server

The Tivoli Storage Productivity Center for Replication server is the control point for the sessions, storage systems, hosts, and other components that are managed by Tivoli Storage Productivity Center for Replication.

### Database

A single database instance serves as the repository for all Tivoli Storage Productivity Center for Replication data.

An embedded database is configured automatically with Tivoli Storage Productivity Center for Replication.

- GUI** The graphical user interface (GUI) is used to manage all Tivoli Storage Productivity Center for Replication functions.
- CLI** The command-line interface (CLI) is used to issue commands for key Tivoli Storage Productivity Center for Replication functions.

## Interfaces for Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication provides a graphical user interface (GUI) and a command-line interface (CLI) for managing data replication and disaster recovery.

### Graphical user interface

Tivoli Storage Productivity Center for Replication contains the following GUI components:

#### Menu bar

Use the menu bar to complete tasks such as navigating the GUI, identifying the ID for the user who is logged on to the GUI, logging out of the GUI, and finding help and other product information.

#### Navigation

The menu bar contains items that open individual pages in the GUI. For example, if you click **Sessions**, the Session page opens. If you click **Storage > Storage Systems**, the Storage Systems page opens. To view the items that are in a menu, hover over the menu.

You can open pages in the GUI by clicking the following items in the menu bar:

#### Overview

Click this item to open the Overview page. This page shows a status summary for all sessions, storage systems, host systems, and management servers that Tivoli Storage Productivity Center for Replication is managing.

The Overview page is displayed when you log on to the GUI.

#### Sessions

Click this item to open the Sessions page. Use this page to create and manage sessions.

#### Storage or Storage > Storage Systems

Click either of these items to open the Storage Systems page. Use this page to create and manage storage system connections and to configure the use of the System Storage DS8000 Easy Tier<sup>®</sup> heat map transfer function.

#### Storage > Host Connections

Click this item to open the Host Connections page. Use this page to create and manage connections to host systems.

#### Storage > Volumes

Click this item to open the Volumes page. Use this page to view the volumes that are associated with a storage system.

### **Paths or Paths > ESS/DS Paths**

Click either of these items to open the ESS/DS Paths page. Use this page to view and manage paths between source logical subsystems and target logical subsystems for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems.

### **Console**

Click this item to open the **Console** window. Use this window to view detailed information about actions that are taken by users, errors that occur during normal operation, and hardware error indications.

### **Settings or Settings > Advanced Tools**

Click either of these items to open the Advanced Tools page. Use this page to create a log file package for diagnostic purposes, change the automatic refresh rate of the GUI, and enable or disable the Metro Mirror heartbeat.




### **Settings > Management Servers**

Click this item to open the Management Servers page. Use this page to manage the active and standby management servers in a high-availability relationship.

### **Settings > Administration**

Click this item to open the Administration page. Use this page to view and manage the users and user groups for Tivoli Storage Productivity Center for Replication and their access privileges.

If one or more of the components that are shown on the correlating GUI page have an Unknown, Severe, or Warning status, one of the following icons is displayed next to the item in the menu bar:

- Unknown status 
- Severe status 
- Warning status 

If there are multiple components on a page and the components have different statuses, the highest level status icon is displayed next to the item in the menu bar. Unknown is the highest level status and Warning is lowest.

### **User ID**

The user ID shows who is logged on to the GUI. An icon next to the user ID indicates the user role. To view the role, hover over the icon.

To log out of the GUI, hover over the user ID and click **Log Out**.

### **Help icon**

The **Help** icon  provides the links to the following items:

- Online help for the active page, window, or notebook in the GUI.
- Production documentation.
- Information about the product.



### Breadcrumb

The breadcrumb shows your location in the GUI. You can click an element in the breadcrumb to jump to the corresponding location in the GUI.

For example, when you view the details page for a session, the following breadcrumb is displayed: **Sessions** > *Session Name*. To return to the Sessions page, click **Sessions** in the breadcrumb.

### Content area

The content area displays a page that correlates to the component that you selected in the menu bar.

### Dynamic images

The GUI provides dynamic images that show the status of the sites and role pairs that are in the session. The images present a pictorial view of information such as the status, progress, and direction of the replication between sites and whether volumes are consistent between sites. Consistent volumes are blue while inconsistent volumes are gray.

Data copying symbols indicate the type of copy that occurs between the sites and the status of the copy. The direction that the symbol is displayed in the GUI depends on the direction of the copy. For example, the following image represents a FlashCopy session. The FlashCopy operation occurs from the H1 volumes to the T1 volumes as indicated by the lightning bolt symbol. The symbol is gray because the FlashCopy operation is inactive and the T1 volumes are gray because they are not consistent.



### Command-line interface

You can use the Tivoli Storage Productivity Center for Replication CLI by using the **csmdi** utilities. You can use the CLI directly or as a script to automate functions.

For security, the CLI runs only on the Tivoli Storage Productivity Center for Replication management server, which is the central point of control for managing data replication. You can run the CLI remotely by using a remote-access utility, for example, Secure Shell (SSH) or Telnet.

For Tivoli Storage Productivity Center for Replication on Windows systems, you can specify remote access to Linux or AIX terminals if you enabled Telnet on your Windows server.

**Related reference:**

“Session images” on page 439

The Tivoli Storage Productivity Center for Replication GUI provides a visual aid to help you create and manage your sessions. The visual aid shows the number of volume roles in the session and how the roles are distributed between the sites. It also shows the copy method and direction.

## Tivoli Storage Productivity Center for Replication integration with Tivoli System Automation Application Manager

IBM Tivoli System Automation Application Manager includes a Distributed Disaster Recovery (DDR) feature that uses IBM Tivoli Storage Productivity Center for Replication to provide a data replication and disaster recovery solution for distributed systems. The DDR feature is available in Tivoli System Automation Application Manager version 3.2 or later and requires Tivoli Storage Productivity Center for Replication version 4.1.1.1 or later.

Tivoli System Automation Application Manager automates the availability of information technology resources. The DDR feature and Tivoli Storage Productivity Center for Replication provide a full end-to-end solution for site switching on distributed systems at the application level.

For more information about the integration of Tivoli Storage Productivity Center for Replication with Tivoli System Automation Application Manager, including the steps required to perform the integration, see the *Tivoli System Automation Application Manager Administrator's and User's Guide*. This guide is provided by version at <http://www.ibm.com/developerworks/servicemanagement/dca/saam/resources.html>.

## Management server

The *management server* is a server that has Tivoli Storage Productivity Center for Replication installed. The management server provides a central point of control for managing data replication.

For more information about the management server, see “Management servers” on page 308.

## Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration.

For more information about the storage systems that you can use with Tivoli Storage Productivity Center for Replication, see “Storage systems” on page 317.

## Host systems

A *host system* is an IBM z/OS or AIX system that connects to storage systems to enable certain replication features for those systems.

A connection to a host system is required in the following scenarios:

- A connection to a z/OS host system is required if you want to enable z/OS features such as HyperSwap<sup>®</sup> and hardened freeze in Tivoli Storage Productivity Center for Replication sessions.
- A connection to an AIX host system is required if you want to use the Open HyperSwap feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site. For the swap to occur, the storage systems that contain the volumes must be attached to an AIX host and the AIX host must be connected to Tivoli Storage Productivity Center for Replication. The Open HyperSwap feature is optional for the Metro Mirror with Failover/Failback session type.

You can add both AIX and z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. For example, you can add a connection to a z/OS host in a Tivoli Storage Productivity Center for Replication instance that is running on Windows, Linux, or AIX.

For more information about the host systems and how to manage them in Tivoli Storage Productivity Center for Replication, see “Managing host systems” on page 333.

## Ports

Tivoli Storage Productivity Center for Replication uses ports for communication with the management servers that are in a high-availability relationship; the graphical user interface (GUI); the command-line interface (CLI); and storage systems.

### Tivoli Storage Productivity Center for Replication

Port numbers are defined during the installation of Tivoli Storage Productivity Center for Replication. To determine the port number for the ports that are used for Tivoli Storage Productivity Center for Replication, refer to the following table.

*Table 15. Where to find Tivoli Storage Productivity Center for Replication port numbers*

Tivoli Storage Productivity Center for Replication component	Location of port information
Web browser	To view the port number for the web browser, open the portdef.props file in the <i>TPC_install_directory\wlp\usr\servers\replicationServer\properties\</i> directory.

Table 15. Where to find Tivoli Storage Productivity Center for Replication port numbers (continued)

Tivoli Storage Productivity Center for Replication component	Location of port information
Standby management server	<p>To view the port number for each management server, click <b>Management Servers</b> in the navigation panel of the Tivoli Storage Productivity Center for Replication GUI. You can also run the <b>lshaservers</b> command from the command-line interface to view the ports for the management servers.</p> <p><b>Important:</b> The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.</p>
Client and CLI port	<p>To view the port number for the client, click <b>About</b> in the navigation panel of the Tivoli Storage Productivity Center for Replication GUI. You can also run the <b>whoami</b> command from the command-line interface to view the port for the client.</p> <p><b>Important:</b> The client port number must be the same on both the management server and the standby management server in a high-availability relationship. If you change the client port number on either the management server or the standby management server, you must also change the port number on the other server.</p>

## Storage system ports

The following table lists the default ports for each storage system type.

Table 16. Storage system default ports

Storage System	Connection Type	Port
<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> </ul>	Direct Connection	2433
<ul style="list-style-type: none"> <li>• System Storage DS8000</li> </ul>	Hardware Management Console Connection	1750
<ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V3500</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>	Direct Connection	443 and 22

Table 16. Storage system default ports (continued)

Storage System	Connection Type	Port
• The XIV system	Direct Connection	7778

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center for Replication can send outgoing TCP/IP packets to the storage controllers.

Because there are typically multiple applications that are running on the management server, it is possible that port conflicts might arise if other applications attempt to use the same ports that Tivoli Storage Productivity Center for Replication uses. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. You can accept the default port or enter another port number.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also, ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

**If the management server is running on Windows Server:** Configure the inbound and outbound rules for Tivoli Storage Productivity Center for Replication. To create a new outbound rule, start the New Outbound Rule Wizard from the Windows Firewall with Advanced Security menu.

## Sessions

A *session* manages a specific type of data replication against a specific set of volumes. The type of data replication that is associated with the session determines the replication actions that are available for the session. For example, the options for FlashCopy sessions are different from the options for Metro Mirror sessions.

For more information about sessions, see “Sessions” on page 340.

## Copy sets

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

For more information about copy sets, see “Copy sets” on page 341.

## Volume roles

A *volume role* is given to every volume in the copy set. The role defines how the volume is used in the copy set and, for multi-site sessions, the site location of the volume. For example, the H1 role is made up of host-attached volumes that are at the primary site.

For more information about volume roles, see “Volume roles” on page 344.

## Role pairs

A *role pair* is the association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is indicated by an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

For more information about role pairs, see “Role pairs” on page 346.

## Practice volumes

You can use a *practice volume* to practice for disaster recovery without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

You can practice disaster recovery actions without using practice volumes. However, without practice volumes, you cannot continue to copy data changes between volumes while you are practicing.

For more information about practice volumes, see “Practice volumes” on page 346.

## Consistency groups

For Global Mirror and Metro Global Mirror sessions, Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that were updated to preserve write order and are recoverable.

For more information about consistency groups, see “Consistency groups” on page 346.

## Failover and failback operations

If a session has *failover* and *failback* capabilities, you can move the operation of Tivoli Storage Productivity Center for Replication from one site to another, and change the direction of the copy without completing a full copy.

Failover is the process of temporarily switching production to a backup facility by following a scheduled maintenance period or a disaster at a production site. A failover operation can be followed by a failback operation, which is the process of returning production to its original location and resuming the copy between the two sites.

For more information about failover and failback operations, see “Failover and failback operations” on page 347.

## Session commands

The commands that are available for a session depend on the session type.

For descriptions of the session commands for each session type, see “Session commands” on page 383.

## Session types

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

The session types that are available depend on the storage system type and the edition of Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication that you are using as shown in this topic.

### Basic HyperSwap

Basic HyperSwap is a copy services solution for IBM z/OS version 1.9 and later. Basic HyperSwap provides high availability of data if a storage system failure occurs.

The Basic HyperSwap session type enables Tivoli Storage Productivity Center for Replication to manage Basic HyperSwap by using Metro Mirror or Metro Global Mirror replication.

Basic HyperSwap sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

*Table 17. Basic HyperSwap*

Supported editions	Supported storage systems
<ul style="list-style-type: none"><li>• Tivoli Storage Productivity Center for Replication Basic Edition for System z</li><li>• Tivoli Storage Productivity Center for Replication for System z</li></ul>	<ul style="list-style-type: none"><li>• TotalStorage Enterprise Storage Server Model 800</li><li>• System Storage DS8000</li><li>• System Storage DS6000</li></ul>

### FlashCopy

FlashCopy replication creates a point-in-time copy of the data on a source volume to a target volume. Data that existed on the target volume is replaced by the copied data.

FlashCopy sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 18. FlashCopy

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication Basic Edition for System z</li> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3500</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Snapshot

Snapshot replication creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Snapshot sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 19. Snapshot

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication Basic Edition for System z</li> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	The XIV system

## Metro Mirror Single Direction

Metro Mirror is synchronous replication that operates between two sites that are up to 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Metro Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are propagated to the target volume before the write operation finishes processing.

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site.

Metro Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.



Table 20. Metro Mirror Single Direction

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Metro Mirror Failover/Failback

Metro Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Metro Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 21. Metro Mirror Failover/Failback

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> <li>• The XIV system</li> </ul>

## Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice replication combines Metro Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 22. Metro Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Global Mirror Single Direction

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Global Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume.

With the Global Mirror Single Direction session type, Global Mirror replication is available only from the primary site.

Global Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 23. Global Mirror Single Direction

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice replication enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Either Direction with Two-Site Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

*Table 24. Global Mirror Either Direction with Two-Site Practice*

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> </ul>

## Global Mirror Failover/Failback

Global Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Global Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

*Table 25. Global Mirror Failover/Failback*

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> <li>• The XIV system</li> </ul>

## Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Global Mirror Failover/Failback with Change Volumes sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 26. Global Mirror Failover/Failback with Change Volumes

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• SAN Volume Controller Version 6.3 and later</li> <li>• Storwize V3700 Version 6.4.1 and later</li> <li>• Storwize V7000 Version 6.3 and later</li> <li>• Storwize V7000 Unified Version 1.3 and later (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice replication combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 27. Global Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Metro Global Mirror

Metro Global Mirror replication provides continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Metro Global Mirror replication enables you to switch the direction of the data flow so that you can use your secondary or tertiary site as your production site.

Metro Global Mirror sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 28. Metro Global Mirror

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800 (only H1 site)</li> <li>• System Storage DS8000</li> </ul>

## Metro Global Mirror with Practice

Metro Global Mirror with Practice replication combines Metro Mirror, Global Mirror, and FlashCopy capabilities to provide a point-in-time copy of the data on the tertiary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Global Mirror with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 29. Metro Global Mirror with Practice

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800 (only H1 site)</li> <li>• System Storage DS8000</li> </ul>

### Related concepts:

“Basic HyperSwap” on page 356

Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

“FlashCopy” on page 356

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

“Snapshot” on page 357

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

“Metro Mirror” on page 358

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is in one storage system and the target is in another storage system.

“Global Mirror” on page 365

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

“Metro Global Mirror” on page 370

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy

into a single session, where the Metro Mirror target is the Global Mirror source.

## Metro Mirror heartbeat

The Metro Mirror *heartbeat* guarantees data consistency across multiple storage systems when the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems.

The Metro Mirror heartbeat is disabled by default. When determining whether to use the Metro Mirror heartbeat, analyze your business needs. If you disable the Metro Mirror heartbeat, data might become inconsistent. If you enable the Metro Mirror heartbeat and a freeze occurs on the storage system, your applications cannot perform write operations during the freeze.

For more information about the Metro Mirror heartbeat, see “Metro Mirror heartbeat” on page 428.

## Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. Site awareness prevents a session relationship from being established in the wrong direction.

For more information about site awareness, see “Site awareness” on page 323.

## Users and groups

For authentication and authorization, Tivoli Storage Productivity Center for Replication uses users and groups.

For information about users and groups for Tivoli Storage Productivity Center for Replication for System z, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

For information about users and groups for Tivoli Storage Productivity Center for Replication, see “Users and groups” on page 478.

## User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

IBM Tivoli Storage Productivity Center for Replication provides the following predefined user roles: monitor, session operator, and administrator.

For more information users roles, see “User roles” on page 479.

---

## Key concepts

This section contains a technical overview that will help you understand how IBM Tivoli Storage Productivity Center works. An understanding of the concepts in this section will help you use Tivoli Storage Productivity Center effectively.

### Data collection

To help you implement a storage management strategy, it is critical that you determine the information that you want to gather about the resources within your environment. You can schedule different data collection jobs depending on the types of information that you want to gather and the resources that you want to monitor.

Use probes, performance monitors, scans, and pings to collect detailed information about resources. You must schedule and run these data collection jobs before you can complete other tasks within Tivoli Storage Productivity Center, such as viewing reports, enforcing storage policies, setting alerts, viewing data paths, and managing storage systems.

You schedule probe and performance monitor jobs for a resource when you add the resource for monitoring in the web-based GUI. After you add the resource, you can modify the probe and performance monitor job schedules on the list and details pages for the resource. For example, to modify the probe job schedule for a storage system, go to the Storage Systems page or the Storage System details page in the web-based GUI.

Use the web-based GUI to schedule data collection jobs to gather the following types of information:

#### **Asset and status information**

Use probes to collect asset, storage statistics, and status information about resources. You can run probes on all the resources that are monitored by Tivoli Storage Productivity Center.

#### **Performance monitoring**

Use performance monitors to collect metrics that measure the performance of switches and storage systems.

Use the stand-alone GUI to schedule data collection jobs to gather the following types of information:

#### **Availability monitoring**

Use pings to collect information about the availability of resources. You can run pings against servers and computers only. Pings are unavailable for storage systems and other resources.

#### **Usage and trending information**

Use scans to collect statistics about the usage and trends of your actual storage consumption. Scans provide all the data for usage reporting, quota analysis, and constraint analysis. They gather information at the file and directory level. You can run scans against servers and computers only. Scans are unavailable for storage systems, switches, fabrics, and other devices.

#### **Tivoli Storage Productivity Center server monitoring**

Use TPC server probes to collect the storage asset and capacity information that was gathered by other Tivoli Storage Productivity Center servers in your environment. Create master and subordinate relationships among

servers to use a single interface to generate reports that are based on data and metrics that are collected by multiple servers.

## Performance monitoring and troubleshooting

Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

To monitor the performance of resources and check for threshold violations, complete the following tasks:

- Add resources for monitoring and schedule data collection
- Define performance thresholds
- View and troubleshoot performance issues

### Collect performance data

Before you can troubleshoot and view reports about performance, you must collect data about monitored resources. Performance monitors are data collection jobs that gather performance information about resources. This information includes metrics that measure the performance of the components within a resource. Metrics measure the performance characteristics of volumes, ports, and disks on storage systems and switches. Tivoli Storage Productivity Center provides many different metrics for measuring performance. For example, some key metrics for storage systems are I/O rate in I/O operations per second, data rate in MiB per second, and response time in milliseconds.

You can use metrics in Tivoli Storage Productivity Center to track growth or change in I/O rates, data rates, and response times. In many environments, I/O and data rates grow over time, and response times increase as those rates increase. This relationship can help with "capacity planning" for your storage. As rates and response times increase, you can use these trends to project when more storage performance and capacity is required.

### Define performance thresholds

Performance thresholds alert you when the performance of a monitored resource falls outside of a specified range and might represent a potential problem. When you define a performance threshold for an internal component of a resource, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, a threshold violation occurs. For example, you can define a performance threshold that alerts you when the overall back-end response time for a managed disk on a SAN Volume Controller exceeds a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

**Restriction:** You only set performance thresholds for certain metrics. For a complete list of performance thresholds, see "Performance thresholds" on page 1100.

### View and troubleshoot performance issues

After data collection and performance thresholds are configured, you can use the web-based GUI to complete the following tasks:



- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance of a resource fell outside of a specific range.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller storage system is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server for monitoring, Tivoli Storage Productivity Center automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, you can view the performance of the internal resources that are directly associated with the SAN storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.
- In the reporting function of the web-based GUI, you can also view and create performance reports about multiple resources.

**Related concepts:**

“Monitoring the performance of resources” on page 205

Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

Chapter 7, “Reporting,” on page 483

Tivoli Storage Productivity Center provides multiple user interfaces for viewing reports about the storage infrastructure in an enterprise environment.

**Related tasks:**

“Adding resources” on page 127

You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

**Related reference:**

“Performance metrics” on page 1069

Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.

“Performance thresholds” on page 1100

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

## Applications and departments

Tivoli Storage Productivity Center provides storage grouping to define resources that utilize application and department business modeling. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. The pages for the monitoring and management of applications and departments are aligned with user scenarios. For example, capacity trending, resource health monitoring and performance troubleshooting.

To manage and administer applications and departments, complete the following tasks:

- Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.
- View the status of resources that make up the applications and departments.
- Create resource filters to automatically add resources to applications.
- Add resources directly to applications
- Add applications to existing applications and departments.
- Add departments to existing departments.
- View subcomponent (member application) information to support hierarchal levels of storage grouping.

### Applications

An application is a program or a project that consumes storage resources within an organization and interfaces with other enterprise groups that are important to the running of a business. Use Tivoli Storage Productivity Center to model the storage usage that is consumed in your environment by assigning the usage to applications to see the overall health status.

An application can be part of a department and have its own subcomponents (member applications) that are used to create a five level deep hierarchy. Applications that are grouped together can range from large line of business systems to specialized software, in a department, that runs on either client computers or servers. For example, an application might be an automated billing system within the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

The Applications page shows capacity information about the applications that are monitored by organization and interfaces with other enterprise groups that are important to the running of a business. Use Tivoli Storage Productivity Center to view and manage the resources that are defined to the applications. If the application is associated with a department, information about that department is also shown.

### Departments

A department is a division within a business. Use Tivoli Storage Productivity Center to model the storage capacity that is consumed in your department for your business environment, in accordance with other department members.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or an application even if they do not belong in the same hierarchy.

For example, in a collaboration scenario, a single IBM SAN Volume Controller might be shared by multiple departments.

The Departments page shows storage capacity information about the top level departments, the subdepartments and any applications that belong to the department that are monitored by Tivoli Storage Productivity Center.

## Storage optimization

Storage optimization tasks optimize the resources in your storage environment. You can use the web-based GUI to create tasks to balance pools, re-tier volumes, or transform volumes that are in storage virtualizer pools.

### Balancing the workload of volumes across storage pools

Use the Balance Pools wizard to balance the workload of volumes across pools on the same tier. The pools are analyzed and recommendations are generated to move volumes from pools with high-activity values to pools with low-activity values.

### Re-tiering volumes

Use the Analyze Tiering wizard to re-tier volumes to balance pools or to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes that are based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

### Transforming volumes

Use the Transform Storage wizard to complete the following tasks for one or more volumes in storage virtualizer pools:

- Move volumes from one storage virtualizer pool to another pool in the same storage virtualizer pool.
- Move volumes in a storage virtualizer pool to a pool that is enabled for Easy Tier.
- Convert fully allocated volumes to thin-provisioned volumes and convert thin-provisioned volumes to fully allocated volumes.
- Convert fully allocated volumes to compressed volumes and convert compressed volumes to fully allocated volumes.

The pools are analyzed and recommendations are generated. For example, analysis occurs to ensure that there is sufficient space in the pool to convert volumes or add volumes.

#### Related tasks:

“Optimizing storage pools” on page 276

Use the Balance Pools wizard to analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Optimizing storage tiering” on page 269

To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

“Managing tasks for tiering storage, balancing pools, and transforming storage” on page 295

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

**Related reference:**

“Criteria for identifying the pools that require balancing” on page 278

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

## Cloud configuration

To take advantage of the simplified provisioning and optimization capabilities that are available by using the web-based GUI, some configuration is required.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, Tivoli Storage Productivity Center can analyze and optimize volume performance. To take advantage of these capabilities, create service classes and, optionally, capacity pools.

A tutorial of cloud configuration concepts is available in the web-based GUI. In the menu bar, go to **Advanced Analytics > Cloud Configuration**, and click **Learn the Concepts**.

### Service classes

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

Service classes simplify provisioning requests by representing a level or type of storage quality. When you are requesting storage, only the required capacity and service class must be specified. Before you can provision storage, you must create service classes that describe the capabilities and characteristics of the storage you want to be able to provision. Service classes can later be modified or deleted as the needs of your installation change.

### Block-storage service classes

A block-storage service class describes attributes and capabilities of block storage resources. When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Tivoli Storage Productivity Center identifies a storage pool for the volume.

### File-storage service classes

A file-storage service class describes attributes of file storage resources. The file storage resources include Network Attached Storage (NAS) filers, Network Shared Disks (NSDs), and file systems. When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, Tivoli Storage Productivity Center identifies a file system or IBM

SONAS NSD for the share. Space on an IBM SONAS NSD is allocated if a new file system is required or an existing file system requires more space.

## Service class properties

A service class represents storage resources with common traits. These common traits are the service class properties, and collectively describe a particular quality of service. Before you create service classes, consider the levels of storage quality that you want to have available for provisioning requests. For block storage, for example, you can organize service classes by storage tier or RAID level. Some service class properties are used during provisioning to identify the best location for storage placement. Other service class properties determine how the storage is configured, and reflect the performance, reliability, or security considerations for the level of service.

In addition to the standard properties of a service class, you can create your own properties by using custom tags. Custom tags can represent any common trait that storage systems must have. For more information about creating your own properties by using custom tags, see “Custom tags” on page 73.

Block-storage service classes can specify the following properties:

- Storage tier or range of storage tiers
- RAID level
- Whether volume virtualization is required or not allowed. This property is available only when at least one IBM storage virtualizer is managed by Tivoli Storage Productivity Center. The IBM storage virtualizers include SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified.  
If volume virtualization is required, the VDisk mirroring property specifies whether volumes are mirrored to a second storage pool, which uses different backend storage, on the same IBM storage virtualizer.
- Whether volumes must be thin provisioned or must not be thin provisioned.  
When a thin-provisioned volume is created in a DS8000, an extent space efficient (ESE) volume is created.  
If thin provisioning is required, you can specify thin-provisioning configuration properties for the IBM storage virtualizers and IBM XIV Storage System. The following properties might be available:
  - If the volume is created on an XIV system, the locking behavior configuration property affects the volume that is created.
  - If the volume is created on an IBM storage virtualizer, the following configuration properties affect the volume that is created:
    - Allocated space
    - Auto-expand
    - Warning level
    - Grain size
- Whether volumes are compressed. This property is applied only when an IBM storage virtualizer has at least one I/O group that contains a compressed volume. If there are no existing compressed volumes in an I/O group, Tivoli Storage Productivity Center will not create compressed volumes.  
This property is available only when the thin provisioning property is enabled.
- A multipathing policy. If a server has a Storage Resource agent (SRA) and uses an IBM System Storage Multipath Subsystem Device Driver (SDD), the server multipathing policy is used to configure the driver. If provisioning is initiated

from the vSphere Web Client extension for Tivoli Storage Productivity Center, the VMware vSphere multipathing policy is used to configure the multipath driver on the hypervisor.

If a multipathing policy is set, and fabrics are managed by Tivoli Storage Productivity Center, the following properties affect how the storage system is connected to the host:

- Whether the storage system and the client host are connected through fully redundant paths
- Number of paths between the storage system and the client host
- If at least one DS8000 Encryption Group is discovered by Tivoli Storage Productivity Center, the following properties are available and affect the storage placement and configuration of volumes:
  - Whether volumes must be encrypted
  - If volumes must be encrypted, the DS8000 encryption group to use

File-storage service classes have the following properties:

- Whether dedicated storage is required, or whether shared storage is allowed. In other words, whether the file system from which the NAS share is provisioned can contain other NAS shares.
- Fileset type. The fileset type is either an independent fileset with its own allocated nodes, or a dependent fileset allocated on the file system.
- Access path host name replacement. For more information, see “Special consideration for provisioning by using the VMware vSphere Web Client extension” on page 74.

## Predefined service classes

To illustrate service classes, a set of predefined service classes are provided. From the Service Classes page in the web-based GUI, you can view any of these service classes for more information. Select a service class and then select **Actions > View / Modify**. To display help information about a service class property, complete the following steps:

1. Move the mouse pointer over the field or control.  
A question mark icon is displayed next to the field.
2. Move the mouse pointer over the question mark icon to display an explanation of the property.

The following sample block-storage service classes are provided for illustration.

**Gold** The Gold service class represents the highest-performing storage resources for mission-critical applications.

**Silver** The Silver service class represents high-performing storage resources for applications in production.

### **Bronze**

The Bronze service class represents standard storage resources for non-mission-critical applications.

The following sample file-storage service classes are provided for illustration.

### **NormalIsolation**

The NormalIsolation service class describes normal isolation file storage by specifying that shared storage is allowed. The file system from which the NAS share is provisioned can contain other NAS shares.

### **EnhancedIsolation**

The EnhancedIsolation service class describes enhanced isolation file storage by specifying that dedicated storage is required. The file system from which the NAS share is provisioned cannot contain other NAS shares.

### **Custom tags**

During provisioning, Tivoli Storage Productivity Center determines candidates for storage placement by comparing the requirements of a service class with the known capabilities, configuration, and performance of the available storage systems. You can think of the storage placement determination as a filtering process. Tivoli Storage Productivity Center filters the set of available storage resources against properties of the service class until only the storage resources that can satisfy all of the requirements remain.

In addition to the standard properties of a service class, you can create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

When you specify tags on a block-storage service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you specify tags on a file-storage service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

During provisioning, Tivoli Storage Productivity Center filters candidates for provisioning against the custom tags.

Before you create service classes, consider whether there are any special requirements for provisioning that are not addressed by the standard properties of the service class. If so, tag the appropriate storage resources to satisfy custom requirements of a service class. When you are defining the service class, specify the custom tags that are required.

### **Adding users to a service class**

By default, only administrators can provision storage. However, an administrator can add users to a service class to grant them permission to provision by using the service class. In order for an administrator to grant permission to provision to a user, the user must be assigned to the Monitor or External Application role.

When you are granting users permission to provision by using a particular service class, you can also specify whether execution of provisioning plans created by these users requires administrator approval. The requirement for administrator approval applies to all non-administrative users who are added to the service class, and also applies to external applications such as the vSphere Web Client extension. Approval cannot be required for some non-administrative users and not for others.

Before you create service classes, consider whether you want to allow users to request their own storage. Consider whether you want user provisioning requests to require administrator approval.

## Associating capacity pools with a service class

Before you create service classes, consider whether certain storage requests must always be satisfied from a particular set of resources. If so, you can associate one or more capacity pools with a service class. For example, suppose that a set of storage resources are allocated to a particular department of your business. All storage requests for that department must be satisfied by those storage resources. In this case, you can add all the storage resources to a capacity pool and associate that capacity pool with a service class. You can also add the members of the department to the service class. The department members can then request their own storage as needed. However, the requests can be satisfied only by the storage resources that are allocated to the department.

When you are configuring capacity pools and service classes, keep in mind that a storage resource can be a member of only one capacity pool. Adding storage resources to a capacity pool is not a viable approach if the resources are already members of another capacity pool. Consider specifying custom tags to achieve the same result. In the preceding example, suppose that some or all of the storage resources that are allocated to the department are already members of capacity pools. In this case, you can instead tag the storage systems and the service class with matching tags to achieve the same result. With either approach, only the storage systems that are allocated to the department are candidates for provisioning when storage is requested by using the service class.

## Special consideration for provisioning by using the VMware vSphere Web Client extension

When a file share is provisioned by the VMware vSphere Web Client extension, an access path to the share is provided by Tivoli Storage Productivity Center. By default, the path includes the cluster name of the file storage system in which the share is created. A cluster name must be a valid host name that is registered with a Domain Name System (DNS) server. If the cluster name is not valid, the share is not accessible under the access path that is returned by Tivoli Storage Productivity Center.

To address this situation, you can set one of its custom tags for a file storage system to specify an access path host name replacement. A file-storage service class can specify which of the file storage system custom tags contains the access path host name replacement.

If file shares are provisioned by using the VMware vSphere Web Client extension, determine whether you must set access path host name replacements.

## Capacity pools

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

The following types of storage resources can be grouped into capacity pools:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM SONAS Network Shared Disks (NSDs)



You can organize storage resources into capacity pools in any way that serves your business needs. For example, you can separate the storage resources that are allocated for separate divisions of your business into separate capacity pools. You can then track the storage use for each division separately, and restrict provisioning requests to the appropriate set of storage resources.

### **Tracking used and available space in a capacity pool**

You can track the used and available space for any set of storage resources by adding the storage resources to a capacity pool. From the Capacity Pools pane of the web-based GUI, you can view the total capacity of the resources, and monitor the following capacity measurements:

- Used file space
- Available file space
- Used block space
- Available block space

### **Restricting a provisioning request to a capacity pool**

You can use capacity pools to define a set of storage resources from which provisioning requests must be satisfied. When you are requesting storage by using the Provision Storage wizard or the vSphere Web Client extension, you can specify a capacity pool. If you do, the provisioning request is restricted to resources in the capacity pool. Only those resources are candidates for provisioning.

You can also associate a service class with capacity pools. If a service class is associated with capacity pools, provisioning requests for the service class must specify, and be constrained to, one of the associated capacity pools. Because you can also grant non-administrative users permission to provision storage by using the service class, associating capacity pools with the service class restricts the users to a specific set of resources.

## **Provisioning storage**

The Provision Storage wizard simplifies the steps for submitting provisioning requests. When you use the Provision Storage wizard, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require.

### **Storage requirements**

In the Provision Storage wizard, you specify your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have in order to provide the service class. For example, when you provision volumes, a block-storage service class specifies properties such as a storage tier, a RAID level, and whether the pool must be able to encrypt or thin provision volumes.

A service class typically represents a quality or service. For example, Tivoli Storage Productivity Center provides predefined block-storage service classes named Gold, Silver, and Bronze. Gold represents the highest-performing storage resources, and Silver and Bronze represent lesser levels of storage quality.

The different levels of storage quality that are provided by each service class are defined by the service class properties. For example, each service class specifies a

different required storage tier that depends on the level of service that is required. Because the Gold service class is for mission critical applications, it sets thin provisioning off.

When you use the Provision Storage wizard, you concern yourself only with which service class is required. If you are unfamiliar with the service classes, however, you can open a separate window to view service class information. If you have Administrator privileges, you can modify or create service classes.

## Storage constraints

Tivoli Storage Productivity Center identifies the storage resources that can provide the capacity and the service class from a set of storage resources. This set might be all the storage resources that are known to Tivoli Storage Productivity Center, or it might be constrained to a subset of those resources. In particular, your site can create capacity pools. *Capacity pools* are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business.

A service class might restrict storage placement to one or more capacity pools. If the service class you are provisioning from restricts placement, you must select one of the allowed capacity pools. If the service class does not restrict storage placement, you can specify that candidates for provisioning can be selected from all storage resources that are known to Tivoli Storage Productivity Center, all storage resources in any capacity pool, or storage resources in a particular capacity pool. If you are unfamiliar with the capacity pools that you can select when you are using the Provision Storage wizard, you can open a separate window to view capacity pool information. If you have Administrator privileges, you can modify or create capacity pools.

If the fabrics are managed by Tivoli Storage Productivity Center, the set of storage resources might be further constrained by your zoning policy. If the zoning policy is not configured to automatic zoning, only storage systems with connectivity to the server or hypervisor are candidates for storage placement.

## Provisioning volumes

**Restriction:** To provision volumes, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.

To provision volumes, you select one or more servers or hypervisors and start the Provision Storage wizard. You can request one or more volumes. For each volume, you specify a volume name, the capacity that is required, a service class, and, optionally, a capacity pool. For each volume, IBM Tivoli Storage Productivity Center identifies the storage pools that can provide the capacity and the service class. From the set of pools that can provide the capacity and service class, Tivoli Storage Productivity Center identifies the best location for the storage. The best location for the storage is based on the unallocated volume space in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is then given to systems that have available performance data.

On storage systems, volumes are allocated in increments of a set unit size that depends on the storage system type. Depending on the storage system that was

identified as the best location for the storage, the volumes that are allocated might be larger than the capacity you requested for the following reasons:

- On SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified, volumes are allocated in increments of an extent size. The extent size is set when the MDisk group is created. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On a DS8000, volumes are allocated in increments of a 1 GiB fixed extent size. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On an XIV system, volumes are allocated in fixed increments of 16 GiB. The capacity that you requested is rounded up, if necessary, to complete the 16 GiB increment.

## Provisioning shares

To provision a share, you select one or more servers or hypervisors and start the Provision Storage wizard. You specify the capacity that is required for the share, a service class, and information about how to export the share. Tivoli Storage Productivity Center identifies the file systems and Network Shared Disks (NSDs) that can provide the capacity and the service class. From the set of file systems and NSDs that can provide the capacity and service class, Tivoli Storage Productivity Center identifies the best location for the storage. The best location for the storage is based on the available space on the file system or NSD.

## The provisioning task

When you complete the steps of the Provision Storage wizard, Tivoli Storage Productivity Center creates a provisioning task for the share or volume. If you request multiple volumes, a separate task is created for each unique service class and capacity pool combination you specify in your volumes request. If a requested volume is not constrained to a capacity pool, the set of all available storage resources is considered the capacity pool. If you have Administrator privileges, you can save, run, or schedule provisioning tasks.

In the service class, an administrator can grant users permission to provision by using the service class. In the service class, an administrator can also specify whether scheduling or running provisioning tasks that are created by using the service class requires administrator approval. If you are a user with permission to provision by using the service class, and administrator approval is not required, you can save, run, or schedule the task. If administrator approval is required, you can only save the task. If you save the task, an administrator can later schedule or run it.

Although Tivoli Storage Productivity Center identifies the best location for storage when it creates the provisioning task, the implementation of the task might fail. For example, the implementation of the task might fail because of changes in the environment that occur after the task is created and before it is run. The implementation of the task might also fail because of storage system restrictions that are not considered when the task is created.

## Viewing task details and logs

When a provisioning task is running, you can view its status in a details page. If you run the provisioning task immediately after you complete the steps of the Provision Storage wizard, the details page is already displayed. You can also

display details of a provisioning task from the Tasks page. From the details page, you can open logs for the provisioning task. The logs show the steps that are taken by the task during processing and include detailed information about any warnings or errors.

## Tracking

In the Provision Storage wizard, you can specify a ticket identifier for tracking purposes. The ticket identifier is associated with the provisioning task, and with any volume or share that is created by the provisioning task.

The ticket identifier can be viewed by showing the Ticket column in the Volumes page or the Shares page. The ticket identifier can also be viewed in the properties notebook for a volume or share.

**Tip:** If you request multiple volumes, a separate provisioning task is created for each unique service class and capacity pool combination you specify in your volumes request. By specifying a ticket identifier when you request the volumes, you can easily track the set of provisioning tasks and volumes that are created by the single request.

## Storage configuration

Whether you are provisioning volumes or shares, Tivoli Storage Productivity Center configures the storage and resources as specified by certain properties of the service class, and, in the case of volumes, according to your zoning policy.

When you provision volumes, Tivoli Storage Productivity Center configures volumes and resources according to properties of the service class:

- If the service class requires thin provisioning, the volume is created in a pool that can thin provision volumes. During provisioning, other properties of the service class are used to configure the thin provisioned volume for the particular type of storage system that contains the pool. Because you configure thin provisioned volumes differently for the different storage system types, the service class has a unique set of thin provisioning properties for each type.
- If the service class specifies a multipathing policy for servers, the IBM System Storage Multipath Subsystem Device Driver (SDD) on the server is configured to use that policy. For the SDD to be configured, a Storage Resource agent (SRA) must be running on the server. If a VMware vSphere multipathing policy is specified for hypervisors, the policy is configured only if provisioning is initiated from the vSphere Web Client extension for Tivoli Storage Productivity Center.

When you are provisioning volumes, if automatic zoning is enabled, new zones might be created to connect a server to the storage system. Existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more new zones are created between a host initiator port and a controller, node, or module port.

When you provision shares, Tivoli Storage Productivity Center configures at least one file access protocol for the file share.

## Tivoli Storage Productivity Center VASA provider

Tivoli Storage Productivity Center VASA provider enables the Tivoli Storage Productivity Center administrator to bring visibility of Tivoli Storage Productivity Center managed storage resources to a vCenter server administrator.

## About vSphere API for Storage Awareness APIs Providers (VASA)

VASA enables coordination between VMware vCenter and storage arrays. It is designed to provide built-in storage insight to support VM storage provisioning, assist with storage troubleshooting and enable new DRS-related use cases for storage.

### VASA benefits

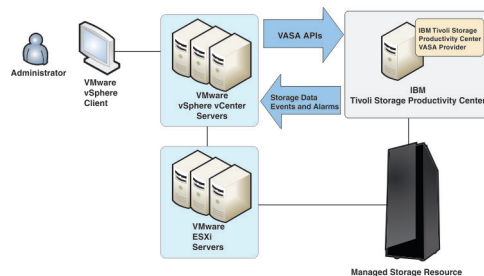
For VMware vSphere users, the Tivoli Storage Productivity Center VASA provider improves the ability to monitor and automate storage-related operations in VMware environments.

A Tivoli Storage Productivity Center storage administrator can share certain Tivoli Storage Productivity Center alerts and storage capabilities for resources that are mapped to a vCenter server.

The Tivoli Storage Productivity Center VASA provider provides the underlying connectivity for alerts and information to be shared.

### Tivoli Storage Productivity Center VASA provider architecture

This figure shows the Tivoli Storage Productivity Center VASA provider running on a Tivoli Storage Productivity Center connected to both the VMware vCenter server and the managed storage resources.



A Tivoli Storage Productivity Center storage administrator can share certain Tivoli Storage Productivity Center alerts for resources that are mapped to a vCenter server.

The Tivoli Storage Productivity Center VASA provider provides the underlying connectivity, and is automatically deployed and running after a Tivoli Storage Productivity Center installation. You must register a Tivoli Storage Productivity Center server to view its information in vCenter storage reports and views.

## vSphere Web Client extension for Tivoli Storage Productivity Center

With the vSphere Web Client extension for Tivoli Storage Productivity Center, you can manage virtual storage in a vSphere environment. You can use the vSphere Web Client extension to complete provisioning tasks and to view information about storage resources in the vSphere Web Client.

## Provisioning storage in the vSphere Web Client extension

You can use the vSphere Web Client extension to provision storage capacity on storage systems that are monitored by Tivoli Storage Productivity Center. You can provision file and block storage and use the service classes and capacity pools that you create in Tivoli Storage Productivity Center. Additionally, you can choose to create data stores on your storage volumes.

## Viewing storage reports in the vSphere Web Client extension

You can use the vSphere Web Client extension to view information about the storage systems that are monitored by Tivoli Storage Productivity Center. You can view the mapping of your virtual storage resources to storage systems that are monitored by Tivoli Storage Productivity Center. You can view metrics for your storage systems. You can also view information about the fabric that is connected to a storage adapter. From the storage reports, you can click links to open the resources in the Tivoli Storage Productivity Center web-based GUI.

## vSphere supported locales

The locales that are supported by the vSphere Web Client extension for Tivoli Storage Productivity Center depend on the vSphere support. In vSphere 5.x, the following locales are supported:

- English: en\_US
- German: de\_DE
- Japanese: ja\_JP
- Simplified Chinese: zh\_CN
- French: fr\_FR
- Korean: ko\_KR

When a locale does not have a translated version, the vSphere Web Client extension is displayed in English, which is the default language.

### Related concepts:

“Provisioning storage with the vSphere Web Client extension” on page 266  
You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage and monitor your virtual environment and storage devices.

“Reporting in a vSphere environment” on page 1013  
You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices.

## Alerts

Determining when and how you are alerted to conditions and violations of performance thresholds within a storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Alerts are triggered by conditions and violations of performance thresholds that are detected during data collection and event processing. For storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system, events are polled every minute from the resource. For the GPFS/GSS storage system, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource

itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.

## Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some triggering conditions, like performance thresholds, require you to enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- A performance threshold was violated
- A data collection job did not complete
- A change occurred in the storage infrastructure

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over utilized or under utilized. Tivoli Storage Productivity Center provides many metrics for measuring performance and determining violations of the thresholds that you specify.

## Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, Tivoli Enterprise Console® events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

**Acknowledging alerts:** Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

## Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable Tivoli Storage Productivity Center to listen to SNMP traps. The default port number is 162, and the default community is public.

**Related concepts:**

“Alerting” on page 159

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines resources for the conditions and performance thresholds that you specify and automatically notifies you when changes or violations are detected.

**Related tasks:**

“Acknowledging alerts” on page 184

Some alerts in Tivoli Storage Productivity Center are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

## Reporting

Use reporting functions to view overview and detailed information about your storage.

From the web-based GUI, you can view predefined reports and create custom reports about the resources managed by Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in IBM Cognos reporting software. Tivoli Common Reporting provides the Cognos software for Tivoli Storage Productivity Center reports.

You can use the Cognos software to view over 70 predefined reports about the capacity and performance of your resources in Tivoli Common Reporting. Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics.

You can schedule reports and specify to create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

In the stand-alone GUI, use the Reporting facilities in the different managers of IBM Tivoli Storage Productivity Center to view both overview and detailed information about your storage. The over 500 available reports are constructed from the statistical information accumulated in the repository and enable you to select exactly the computers, storage systems, switches, clusters, disks, file systems, hypervisors, users and even other Tivoli Storage Productivity Center servers upon which you want to report. You can view reports as tables of data or as graphical charts, which provides you with the flexibility to get a very detailed look at your information or a very high level overview.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use monitoring jobs in Tivoli Storage Productivity Center such as probes, scans, pings, and Tivoli Storage Productivity Center server probes, and discovery jobs to gather comprehensive information and statistics about your storage resources.

## Units of measurement for storage data

Tivoli Storage Productivity Center uses decimal and binary units of measurement to express the size of storage data.



Decimal units such as kilobyte (KB), megabyte (MB), and gigabyte (GB) are commonly used to express the size of data. Binary units of measurement include kibibyte (KiB), mebibyte (MiB), and gibibyte (GiB). Table 30 compares the names, symbols, and values of decimal and binary units.

Table 30. Comparison of binary and decimal units and values

Binary			Decimal		
Name	Symbol	Value (base 2)	Name	Symbol	Value (base 10)
kibibyte	KiB	2 <sup>10</sup>	kilobyte	KB	10 <sup>3</sup>
mebibyte	MiB	2 <sup>20</sup>	megabyte	MB	10 <sup>6</sup>
gibibyte	GiB	2 <sup>30</sup>	gigabyte	GB	10 <sup>9</sup>
tebibyte	TiB	2 <sup>40</sup>	terabyte	TB	10 <sup>12</sup>
pebibyte	PiB	2 <sup>50</sup>	petabyte	PB	10 <sup>15</sup>
exbibyte	EiB	2 <sup>60</sup>	exabyte	EB	10 <sup>18</sup>

Binary units of measurement express the size of data more accurately. When you compare the size of 100 KB to 100 KiB, the difference is relatively small, 2.35%. However, this difference grows as the size of the data values increases. When you compare the size of 100 TB to 100 TiB, the difference is 9.06%.

In general, Tivoli Storage Productivity Center uses base 2 values for memory and disk space values, and base 10 values for space on physical hard drives.

Table 31 shows the percentage difference between decimal and binary values across a range of data sizes.

Table 31. Percentage difference between decimal and binary units

Decimal value	Binary equivalent of decimal value	Difference
100 kilobytes (KB)	97.65 kibibytes (KiB)	2.35%
100 megabytes (MB)	95.36 mebibytes (MiB)	4.64%
100 gigabytes (GB)	93.13 gibibytes (GiB)	6.87%
100 terabytes (TB)	90.94 tebibytes (TiB)	9.06%
100 petabytes (PB)	88.81 pebibytes (PiB)	11.19%
100 exabytes (EB)	86.73 exbibytes (EiB)	13.27%

## Role-based authorization

Roles determine the functions that are available to users of Tivoli Storage Productivity Center. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

In Version 5.2, the roles that were previously defined in Tivoli Storage Productivity Center were consolidated into a smaller set of roles. The following table shows how the roles in versions earlier than 5.2 are mapped to the current set of roles:

Table 32. How roles in previous versions of Tivoli Storage Productivity Center are mapped to the roles in version 5.2 and later

Roles in previous versions	Roles in version 5.2 and later	Authorization level
Superuser Productivity Center administrator Disk administrator Fabric administrator Data administrator Tape administrator	Administrator	This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role. <b>Note:</b> When Tivoli Storage Productivity Center is first installed, the following operating system groups are assigned the Administrator role: <ul style="list-style-type: none"> <li>• Windows: Administrators</li> <li>• UNIX and Linux: root</li> <li>• AIX: system</li> </ul>
Disk operator Fabric operator Data operator Tape operator	Monitor	In the web-based GUI, this role has access to the following read-only functions: <ul style="list-style-type: none"> <li>• Viewing and exporting information about monitored resources</li> <li>• Viewing and acknowledging alerts</li> <li>• Viewing tasks and data collection jobs</li> <li>• Viewing data paths</li> <li>• Opening management GUIs</li> <li>• Opening logs</li> <li>• Tiering storage</li> </ul> In the stand-alone GUI, this role has access to the following functions: <ul style="list-style-type: none"> <li>• Viewing data that is collected by Tivoli Storage Productivity Center</li> <li>• Creating, generating, and saving reports</li> </ul>
This role did not exist in versions 5.1 or earlier.	External Application	This role enables users of other applications to use the provisioning capability of Tivoli Storage Productivity Center to provision storage. For example, a VMware user with this role can provision storage in the vSphere GUI by using the vSphere Web Client extension for Tivoli Storage Productivity Center.  If you assign the External Application role to the user, you must also assign one or more service classes to the user.  This role does not enable users to log in to the Tivoli Storage Productivity Center GUIs.

**Tips:**

- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a

user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.

- If a user is not a member of a group that is assigned a Tivoli Storage Productivity Center role, no access is granted to that user.
- The role of the user who is logged in to the web-based GUI is displayed in parenthesis in the upper-right corner of the banner at the top of every page.
- For rollup reports, you need Administrator authority to do the following actions:
  - Add, remove, or modify the Tivoli Storage Productivity Center subordinate server that the master server is monitoring.
  - Create or run probe jobs that include Tivoli Storage Productivity Center subordinate servers.

Any Tivoli Storage Productivity Center role can generate rollup reports.

- The Tivoli Storage Productivity Center installation program adds the administrator, external application, and monitor roles to the Tivoli Storage Productivity Center installation user.

## Fabrics and zones

Use the Tivoli Storage Productivity Center to learn more about fabrics and zones.

### Fabrics and zones in the web-based GUI

You can monitor switches and fabrics in the web-based GUI. You can also enable automatic zoning so that new zones are created during storage provisioning, if necessary, to connect a server to a storage system.

In the web-based GUI, you can add switches and fabrics for monitoring. You can add the following types of switches:

- Brocade
- Cisco
- QLogic

After switches and fabrics are probed, you can view detailed status and performance information about the switches and fabrics. You can also test the connection to the switch or fabric.

Information that you can view about a fabric includes the following properties:

- The fabric type
- The number of physical and virtual switches in the fabric
- The number of ports that are on all the switches in the fabric
- The number of switch ports that are online and connected to other ports in a fabric
- The name of the principal switch of the fabric

Information that you can view about a switch includes the following properties:

- The number of ports on a switch
- The number of ports that are connected to a storage resource
- The name of the fabric where a switch is a member
- The name of the vendor or manufacturer for a switch

From the Details page of a fabric, you can open the Zone Sets page to view information about the zone sets on the fabric, and to identify which is the active zone set.

You can enable automatic zoning in the web-based GUI. When automatic zoning is enabled, Tivoli Storage Productivity Center creates new zones during provisioning if new zones are needed to connect the storage system with the server. When you set the zoning policy, you can specify whether zone changes from automatic zoning are made to the active zone set, or to a new inactive active zone set.

## Zones, zone aliases, and zone sets

Zones, zone aliases, and zone sets allow logical grouping of ports and storage devices within a storage area network. This section describes zoning concepts and elements.

In a storage area network a *zone* is a logical grouping of ports to form a virtual private storage network. Zones that belong to a single SAN can be grouped into a *zone set*, which can be activated or deactivated as a single entity across all switches in the fabric. A zone set can contain one or more zones, and a zone can be a member of more than one zone set. Using zoning, you can automatically or dynamically arrange fabric-connected devices into logical groups across a physical fabric.

Ports and devices in a zone are called *zone members*. A zone can contain one or more zone members. Ports that are members of a zone can communicate with each other, but they are isolated from ports in other zones. Devices, however, can belong to more than one zone. A *zone alias* is a collection of zone members. A zone alias can be added to one or more zones.

**Note:** In some cases, inactive zones might not have any zone members. Activating an inactive zone set that contains empty zones will fail if a switch does not support empty zones in active zone definitions.

Zone membership can be specified by:

- The N\_Port\_Name of the N\_Port connected to the switch (also known as WWN zoning or port name zoning)
- The N\_Port address identifier assigned during fabric login (also known as PortId or FCID zoning)
- The Node\_Name associated with the N\_Port
- The Domain identification (Domain\_ID) and physical port ID of the Switch Port to which the N\_Port is attached (also known as domain port zoning)
- An alias name

Zoning supports the use of aliases, which are meaningful names assigned to devices. An alias can also be a group of devices that are managed together to make zoning easier.

There are two types of zoning:

### Hardware zoning (port zoning)

In hardware zoning (also called *port zoning*), the members of a zone are the physical ports on a fabric switch.

### Software zoning (WWN zoning)

Software zoning uses the Simple Name Server (SNS) that runs inside a

fabric switch. It is based on the node WWN or port WWN of the zone members to be included. Software zoning lets you create symbolic names for the zones and zone members.

A *default zone* is a group of devices that are not members of the active zone set. These can communicate with each other but not with members of any other zone. Default zoning is enabled by default. You can use a switch element manager to configure the Default Zone option to enable or disable the default zone independently of the active zone set.

**Note:**

1. If the default zone is disabled, devices that are not members of the active zone set cannot communicate.
2. If the default zone is disabled and no zone set is active, no devices can communicate.
3. If default zoning is enabled, deactivating the active zone set makes all devices members of the default zone. If default zoning is disabled, all communication stops.
4. If you activate one zone set while another zone set is active, the currently active zone set is deactivated.
5. If your EFC Manager manages multiple fabrics, ensure that you have the correct zone set for the fabric you are currently updating.

You can view zones in SANs using the topology viewer console. Fabric Manager lets you view zone membership.

**Note:** Certain types of switches and the types of agents configured can affect the zoning operations that you can perform.

### **Switch zoning capabilities**

When you select a fabric in which to perform zoning operations, Fabric Manager determines the capabilities of the switches in the SAN and limits the zoning operations based on that information.

To see a list of supported switches and their capabilities, default values, ranges, and possible effects, go to <http://www.ibm.com/support/docview.wss?uid=swg27027300>.

### **Zoning for Cisco MDS 9000 devices:**

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. Fabric Manager provides basic zone discovery for the Cisco MDS 9000 series of devices using the FC-GS-3 protocol.

The following tasks are supported:

- Creating, deleting, and updating zones
- Creating, deleting, updating, activating, and deactivating zone sets

Fabric Manager supports these zone member types:

- N\_Port WWN
- FC ID (the fibre channel ID of an N\_port attached to the switch)

Unsupported zone and zone member types are displayed as non-standard on the topology display. Fabric Manager does not allow zone management of zones and zone sets that contain unsupported members.

#### **Brocade switches in zones:**

Brocade switches support some non-standard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, Fabric Manager preserves them and does not modify them in any way. If they are part of a zone set that is active at some time, the devices that are part of such zones that are also online are displayed in the topology Zone View.

You can create, change, and delete non-standard zones by using the Brocade switch management application.

**Note:** Tivoli Storage Productivity Center allows you to add empty zones to an inactive zone definition. Activation of a zone set containing empty zones may still fail if the switch does not support empty zones in active zone definitions.

## **Agents**

The IBM Tivoli Storage Productivity Center uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, Fabric agents, and out-of-band Fabric agents.

### **Common Information Model agents**

Common Information Model (CIM) agents enable communication between storage devices, such as storage systems and fabric switches, and IBM Tivoli Storage Productivity Center. Each storage device that you want to manage must have a CIM agent either installed or embedded unless it is a device that uses Native API connection.

CIM agents are provided by the vendor of the storage subsystem or fabric switch. Each vendor provides unique agent code for their family of storage devices. This code implements a Common Information Model Object Manager (CIMOM) that conforms to the Storage Management Initiative Specification (SMI-S) of the Storage Networking Industry Association (SNIA).

The CIM agent enables communication between the storage device and Tivoli Storage Productivity Center. Commands and responses are transmitted between Tivoli Storage Productivity Center and the CIM agent using an XML transport layer. The CIM agent to storage device layer uses a vendor-specific proprietary interface.

The CIM agent usually must be installed and configured, so that it can identify the storage devices with which it communicates. Some storage devices, such as Cisco fabric switches, contain embedded CIM agents and so do not require that CIM agents be installed, although Tivoli Storage Productivity Center must be configured to point directly to the storage devices that contain the embedded CIM agents.

CIM agents can be referred to as SMI-S providers, CIM proxy agents, and CIMOM agents. CIM agents can be embedded in the device or installed on a separate computer.

**Note:**

- Do not install multiple CIM agents on a single computer because of port conflicts.
- Do not install a CIM agent on the system where a Tivoli Storage Productivity Center server component is installed.

### **Storage Resource agents**

Use Storage Resource agents to collect information about storage resources such as servers, virtual machines, workstations, HBAs, and fabrics.

You must deploy Storage Resource agents on resources where you want to gather the following information:

- Asset information
- File and file system attributes
- Database application information
- Network-attached storage (NAS) device information
- Topology information
- Information about zoning and the fabrics that are visible to the server

**Tip:** You can also monitor servers without deploying a Storage Resource agent. When you add an agentless server, Tivoli Storage Productivity Center correlates information about that server with the known host connections on monitored resources. If a match is made between the server and a monitored resource, you can view topology information and the capacity and performance of the storage that is assigned to that server. For more information about when to deploy a Storage Resource agent or add an agentless server, see “Adding servers” on page 138.

### **Out-of-band fabric agents**

An out-of-band agent is software that runs on a switch. Tivoli Storage Productivity Center uses SNMP to communicate with the out-of-band agent to retrieve information about the switch and the fabric.

Be sure that your out-of-band fabric agents point to each switch in the Storage Area Network (SAN) that you want to manage.

Out-of-band fabric agents are necessary for Virtual storage area network (VSAN) information for Cisco switches. For Brocade fabrics, if no SMI-S agent is configured, Out-of-band fabric agents are necessary for Zoning information for Brocade switches.

Tivoli Storage Productivity Center uses SNMPv1 to probe switches and fabrics, and uses SNMPv2 to collect performance data. Tivoli Storage Productivity Center does not support SNMPv3.

For more information about planning for out-of-band fabric agents and Cisco switches, see the *Installation and Configuration Guide*.

### **Tivoli Storage Productivity Center Monitoring Agent**

The IBM Tivoli Storage Productivity Center Monitoring Agent can be used by the IBM Tivoli Enterprise Monitoring Server to monitor systems in your enterprise. This agent is an optional program you can install and use in your enterprise.

IBM Tivoli Monitoring monitors and manages system and network applications on a variety of operating systems, tracks the availability and performance of your

enterprise system, and provides reports to track trends and troubleshoot problems. The Tivoli Storage Productivity Center Monitoring Agent participates in that environment.

For performance reasons, install the Tivoli Storage Productivity Center server on a separate system from the IBM Tivoli Monitoring server. You can install the Tivoli Storage Productivity Center Monitoring agent on any system you want to monitor.

Before using the Tivoli Storage Productivity Center Monitoring Agent, you must have IBM Tivoli Enterprise Monitoring Server installed. For information about the IBM Tivoli Enterprise Monitoring Server, see [http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc\\_6.2.1/welcome.htm](http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.1/welcome.htm).

For information about installing IBM Tivoli Monitoring, see the Quick Start Guide at <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc/GI11-8058-00.htm>.

To use the Tivoli Storage Productivity Center Monitoring Agent, install this agent on the system you want to monitor.

To view the type of reports that are displayed by the Tivoli Storage Productivity Center Monitoring Agent, see [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.1/com.ibm.tpc\\_V521.doc/fqz0\\_c\\_using\\_tpcma.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.1/com.ibm.tpc_V521.doc/fqz0_c_using_tpcma.html).

## Single sign-on

You can use the *Single sign-on* authentication process to enter one user ID and password to access multiple applications. *Single sign-on* integrates with the launch and launch-in-context features to enable the smooth movement from one application to another application.

For the current version of Tivoli Storage Productivity Center, *single sign-on* is available for the following applications:

- Tivoli Storage Productivity Center for Replication
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- Storwize V3500
- Storwize V3700
- IBM Storwize V7000

For Tivoli Storage Productivity Center for Replication, *single sign-on* is enabled by default.

For DS8000 Storage Manager, SAN Volume Controller, and Storwize V7000, *single sign-on* must be enabled.

To enable *single sign-on*, see the following topics in the Tivoli Storage Productivity Center documentation:

- Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication
- Configuring Tivoli Storage Productivity Center for SAN Volume Controller and Storwize V7000 *single sign-on*



## How single sign-on works

A *single sign-on* environment requires a centralized authentication repository that is accessed by all applications within the environment. When you log on to Tivoli Storage Productivity Center, a Lightweight Third-Party Authentication (LTPA) token is created and used for *single sign-on* authentication when you start other applications from within the Tivoli Storage Productivity Center GUI. LTPA is the security technology that is used in the IBM WebSphere Application Server for passing the user authentication information between applications. To use the LTPA tokens, each application must have the same set of public keys to encrypt the user's information. The authenticating service uses the corresponding private keys to decrypt the user's information for authentication. As an extra security mechanism, LTPA tokens expire after a determined amount of time. By default, the tokens expire after 8 hours. You can change the LTPA token expiration time that uses the IBM WebSphere Application Server administration console.

During the period between when you log on to Tivoli Storage Productivity Center and when you start other applications, such as element managers, from the Tivoli Storage Productivity Center GUI, the following conditions might occur:

Condition	Explanation
The user password that was used to log on to Tivoli Storage Productivity Center is changed in the user repository.	Under this condition, the original user credentials that were used to log in to Tivoli Storage Productivity Center are used to access other applications until the timeout period for the LTPA token that is used for <i>single sign-on</i> expires. When the LTPA token expires, you are prompted to reenter your user ID and password when you attempt to start another application that uses <i>single sign-on</i> .
The user repository is not accessible.	Under this condition, the <i>single sign-on</i> feature does not work. You are always prompted to reenter your user ID and password when you attempt to start another application through <i>single sign-on</i> .

Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication use their respective IBM WebSphere Application Server instances to authenticate LTPA tokens. However, other applications, such as the System Storage DS8000 element manager and other element managers that do not run within an IBM WebSphere Application Server instance must be configured for LDAP.

The *single sign-on* feature is not supported by the Tivoli Storage Productivity Center command-line interface (CLI).

## Selecting the user authentication method to use with single sign-on

During the installation, Tivoli Storage Productivity Center uses the operating system (OS) users and groups on the server where Tivoli Storage Productivity Center is installed for authentication and authorization. After the installation, you can configure Tivoli Storage Productivity Center to use LDAP as the authentication and authorization repository.

### OS Authentication

This method authenticates the user against the users that are defined for the local operating system (OS).

### LDAP/Active Directory

This method authenticates the user against a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository.

If OS authentication is selected, the use of the *single sign-on* feature is limited. OS authentication does not support *single sign-on* for element managers, even when the element manager is installed on the same computer as Tivoli Storage Productivity Center. LDAP or Microsoft Active Directory authentication supports *single sign-on* for element managers regardless of where they are installed.

You can change the user authentication method you use for the WebSphere Integrated Solutions Console that is associated with the Tivoli Storage Productivity Center 5.2 (or later) Web Server

## Launch and launch in context

The launch and launch in context features are used to start the Tivoli Storage Productivity Center stand-alone GUI and to start other applications from Tivoli Storage Productivity Center. The terms inbound and outbound are used to differentiate starting Tivoli Storage Productivity Center (inbound) and starting another application from Tivoli Storage Productivity Center (outbound).

Inbound launch in context refers to starting Tivoli Storage Productivity Center at a specific point in the stand-alone GUI. You can use inbound launch in context in either of the following ways:

- Configure an application such as Tivoli Enterprise Portal to start Tivoli Storage Productivity Center at functionally related locations in the stand-alone GUI.
- Start Tivoli Storage Productivity Center from Java Web Start, the command line, or the **Start** menu or Start page in Windows. Use parameters to specify the starting location in the stand-alone GUI.

For more information about inbound launch for the stand-alone GUI, see Inbound launch in context.

Outbound launch refers to starting an application at the starting point for that application. Use outbound launch to open the management GUIs of monitored storage systems and switches from the Tivoli Storage Productivity Center web-based GUI. For more information about outbound launch in the web-based GUI, see “Opening the management GUI for a resource” on page 284.

---

## Chapter 4. Scenarios

This section contains scenarios which illustrate how to perform specific tasks using IBM Tivoli Storage Productivity Center.

The provided scenarios set out situations that can be solved using Tivoli Storage Productivity Center. When using these scenarios you should remember the following:

- The scenarios are intended to be step-by-step instructions that you follow to complete a specific job. Each scenario contains multiple tasks that must be followed in the order given.
- Specific values are given in the scenarios, such as user name, IP Address, probe name, and so forth. These are for illustration purposes only and you must replace them with values appropriate for your system.
- The scenarios do not provide in-depth conceptual information about the tasks. See the appropriate topics in the IBM Tivoli Storage Productivity Center Information Center for more information about the concepts behind the tasks.
- You will complete a scenario successfully if you follow the instructions correctly and your system performs as expected.

---

### Optimizing the performance of storage virtualizers

In these scenarios, you run the Balance Pools and the Analyze Tiering wizards to help resolve performance hot spots and manage storage resources efficiently.

You can optimize the placement of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

To optimize the placement of volumes in storage virtualizer pools, you must complete the following tasks:

- Probe the storage virtualizers.
- Set the tier level of the storage pools that you want to analyze.
- Collect performance data to gauge the utilization percentage for pools and the average workload activity of the volumes.

With the Balance Pools wizard, you optimize storage performance by redistributing volume workloads across pools on the same tier.

With the Analyze Tiering wizard, you optimize storage performance by re-tiering volumes to balance pools and by setting minimum thresholds to regulate the tiering of volumes.

#### **Related tasks:**

“Optimizing storage tiering” on page 269

To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

“Optimizing storage pools” on page 276

Use the Balance Pools wizard to analyze the activity of pools and resolve

performance hot spots by redistributing volumes across each storage tier.

“Managing tasks for tiering storage, balancing pools, and transforming storage” on page 295

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

## Redistributing volume workloads to balance pools

In this scenario, you identify performance hot spots and then run the Balance Pools wizard to help resolve them.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is `svc_payment_records`, and the names of the DS8000 back-end storage systems are `ds_pr1` and `ds_pr2`.

If the activity deviation percentage of the pool exceeds the threshold set for the pool, which is  $> 10\%$ , the pool is identified as a candidate for balancing. If the activity deviation percentage of the pool lies within the threshold range set for the pool, which is  $< -10\%$  to  $> +10\%$ , the pool is considered to be balanced.

**Tip:** On the Pools page, the activity level of the pool is shown in the **Activity** column. The difference between the activity level of the pool and all of the other pools on the same tier and storage system is shown in the **Activity Deviation (%)** column.

The information in Table 33 is provided to illustrate how the volumes are balanced in the pools on the storage virtualizer.

*Table 33. Activity deviation percentages for the pools on the storage virtualizer `svc_payment_records`*

Back-end storage system	Tier	Pool	Activity deviation percentage (%)
<code>ds_pr1</code>	1	<code>pool_ds81_1</code>	20%
<code>ds_pr2</code>	1	<code>pool_ds82_101</code>	-20%
<code>ds_pr1</code>	2	<code>pool_ds81_2</code>	20%
<code>ds_pr1</code>	2	<code>pool_ds81_3</code>	-30%
<code>ds_pr2</code>	2	<code>pool_ds82_103</code>	12%
<code>ds_pr1</code>	3	<code>pool_ds81_4</code>	5%
<code>ds_pr1</code>	3	<code>pool_ds81_5</code>	-3%
<code>ds_pr2</code>	3	<code>pool_ds82_104</code>	0%

When you examine the activity deviation percentage of the pools on the Pools page of the web-based GUI, you make the following decisions:

- To redistribute volumes across the pools on tier 1 because the activity percentage of the `pool_ds81_1` pool is 10% higher than the threshold level.

- To redistribute volumes across the pools on tier 2 because the pool\_ds81\_2 pool is 10 % higher than the threshold level, and the pool\_ds82\_103 pool is 2 % higher than the threshold level.

The pools on tier 3 do not require balancing.

To balance the pools, you open the Balance Pools wizard.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. To group the pools by storage system, click the **Storage System** column heading. Examine the activity deviation percentage values for the pools on the **svc\_payment\_records** storage virtualizer.
3. Select the pools on tier 1 and tier 2, right-click, and then select **Balance Pools**.
4. To specify the number of days for estimating the activity of the pools, type 5.
5. Click **Advanced balancing** to set an activity limit for the pools on tier 1 and tier 2. To ensure that the performance of the pools does not deteriorate when volumes are added, you set an activity limit for all of the pools on each tier.
6. Click **Analyze**.

The Balance Analysis task opens.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule > Analysis**.

**Related tasks:**

“Optimizing storage pools” on page 276

Use the Balance Pools wizard to analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Modifying the criteria for balancing pools” on page 278

When a task is created after you run the Balance Pools wizard, you can modify the criteria that determine which pools are balanced.

**Related reference:**

“Criteria for identifying the pools that require balancing” on page 278

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

## Re-tiering volumes to balance pools

In this scenario, you must re-tier some of the volumes to help resolve the performance hot spots that you identified. To balance and re-tier volumes, you run the Analyze Tiering wizard.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is **svc\_sales\_records**, and the names of the two DS8000 back-end storage systems are **ds81** and **ds82**.

If the activity deviation percentage of the pool exceeds the threshold of > 10 %, the pool is identified as a candidate for balancing. If the activity deviation percentage of the pool lies within the threshold range set for the pools, which is < -10 % to > +10 %, the pool is considered to be balanced.

The information in Table 34 on page 96 is provided to illustrate which pools require the re-tiering of volumes to balance the pools.

Table 34. Activity deviation percentages for the pools on the storage virtualizer  
 svc\_sales\_records storage virtualizer

Back-end storage system	Tier	Pool	Activity deviation percentage (%)
ds81	1	pool_ds81_1	11%
ds82	1	pool_ds82_101	-12%
ds81	2	pool_ds81_2	13%
ds81	2	pool_ds81_3	12%
ds82	3	pool_ds81_5	5%
ds82	3	pool_ds82_104	0%

In the following scenario, the pools are balanced by moving volumes to pools on the same storage tier and by moving volumes to higher storage tiers.

- On tier 3, the activity deviation percentage of all of the pools lies within the threshold range of < -10 % to > +10 %. The pools on storage tier 3 do not require balancing; therefore, no recommendations are generated.
- On tier 2, the activity deviation percentage of all of the pools exceeds the threshold value of 10%. Because the volumes cannot be redistributed to pools on the same storage tier, recommendations are generated to move the most active volumes from the pools on storage tier 2 to pool\_ds82\_101 on storage tier 1.
- On tier 1, the activity deviation percentage of the pool\_ds81\_1 pool, exceeds the threshold value of 10%. Because volumes from storage tier 2 were added to the pool\_ds82\_101 pool, the current activity deviation percentage is no longer accurate. However, because the activity level of the pool\_ds82\_101 pool is below average, recommendations are generated to move volumes on storage tier 1 from the pool\_ds81\_1 pool to the pool\_ds82\_101 pool.

To help resolve the performance hot spots, you open the Analyze Tiering wizard.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Right-click **svc\_sales\_records**, and then select **Analyze Tiering**. Because you open the Analyze Tiering wizard at the storage system level, all of the pools on the storage virtualizer are selected for analysis.
3. Select all of the pools as target pools, and then click **Next**. The pools that you selected for analysis are now also selected as target pools.
4. On the Optimize the Placement of Volumes page, complete these steps:
  - a. To specify the number of days for estimating the activity of the pools, type 5.
  - b. Set an activity limit for the pools for each tier. To ensure that the performance of the pools does not deteriorate when volumes are added, you set an activity limit for all of the pools on each tier.
5. Click **Analyze**.

The Analyze Tiering task is shown.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule > Analysis**.

**Related tasks:**

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

**Related reference:**

“Criteria for identifying the pools that require balancing” on page 278  
To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

## Tiering volumes by I/O density and file usage

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To ensure that only volumes that require the performance capability of tier 1 storage, you add minimum thresholds for I/O density and file usage for each storage tier.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is `svc_sales_records`, and the names of the DS8000 back-end storage systems are `ds81` and `ds82`. The pools on the storage virtualizer are assigned to tier 1, tier 2, and tier 3.

With the Analyze Tiering wizard, you want to analyze the placement of the volumes on tier 1. You set minimum threshold values for I/O density and file usage to re-tier the volumes that do not meet the minimum threshold criteria that you set for I/O density and file usage.

On the Optimize the Placement of Volumes page of the Analyze Tiering wizard, you want to set the following minimum thresholds for each tier.

*Table 35. Minimum threshold values for tier 1, tier 2, and tier 3*

Tier	I/O density minimum threshold level	File usage minimum threshold level
1	4 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 1 day.
2	3 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 1 week.
3	1 operation per second per gigabyte of storage	Eighty percent of files were last accessed within 1 month.

The information in Table 36 is provided to illustrate the tiering of volumes.

*Table 36. Current<sup>®</sup> I/O density and file usage data for the volumes in tier 1 pools*

Pool	Volumes	Current I/O density	Current file usage
pool_ds81_1	volume_1	4 operations per second per gigabyte of storage	Ninety percent of the files were last accessed within one day.
pool_ds81_1	volume_2	2 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 3 weeks.
pool_ds82_101	volume_11	2 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 2 weeks.

Table 36. Current<sup>®</sup> I/O density and file usage data for the volumes in tier 1 pools (continued)

Pool	Volumes	Current I/O density	Current file usage
pool_ds82_101	volume_12	3 operation per second per gigabyte of storage	Eighty-five percent of files were last accessed within 1 week
pool_ds82_101	volume_13	5 operations per second per gigabyte of storage	Ninety-five percent of the files were last accessed within one day.

When you run the Analyze Tiering wizard, the I/O density and file usage data of the volumes on tier 1 is compared with the minimum threshold levels that you set for the storage tiers. The volumes in the tier 1 pools with values lower than the minimum threshold levels are selected as candidate volumes for re-tiering. The criteria that determines the destination pool for the candidate volumes are as follows:

- The minimum threshold levels that you set for tier 2 and tier 3 storage on the Optimize the Placement of Volumes page of the Analyze Tiering wizard.
- The activity limits that you set for destination pools on the Optimize the Placement of Volumes page of the Analyze Tiering wizard.

In this scenario, the following recommendations are generated:

- Move volume\_12 from tier 1 to a pool on tier 2.
- Move volume\_2 and volume\_11 from tier 1 to a pool on tier 3.

Because volume\_1 and volume\_13 meet the minimum threshold levels that are set for I/O density and file usage, no recommendations are generated to move the volumes.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Double-click **svc\_sales\_records**.
3. In the Internal Resources section, double-click **Pools**, and then click the **Tier** column heading to group the pools by tier.
4. Select all of the pools on tier 1, then right-click and select **Analyze Tiering**.
5. Select all of the pools on tier 2 and tier 3 as target pools, and then click **Next**. To group the pools by tier, click the **Tier** column heading.
6. On the Optimize the Placement of Volumes page, specify the number of days for estimating the activity of the pools.
7. Select **Show advanced options**.
8. To specify the activity limit and minimum threshold levels for placing volumes on tier 1, complete the following steps:
  - a. Set an activity limit for the tier. To ensure that the performance of the pools does not deteriorate when volumes are added, you set an activity limit for all of the pools on each tier.
  - b. Select **I/O density per second per gigabyte**, and type 4.
  - c. Select **Percentage of files**, type 80 for the percentage of files, and ensure that the period is set to 1 day.
9. To specify the minimum thresholds for placing volumes on tier 2, complete the following steps:
  - a. Set an activity limit for the tier.



- b. Select **I/O density per second per gigabyte**, and type 3.
  - c. Select **Percentage of files**, type 80 for the percentage of files, and ensure that the period is set to 1 week.
10. To specify the minimum thresholds for placing volumes on tier 3, complete the following steps:
- a. Set an activity limit for the tier.
  - b. Select **I/O density per second per gigabyte**, and type 1.
  - c. Select **Percentage of files**, type 80 for the percentage of files, and ensure that the period is set to 1 month.
11. Click **Analyze**.

The Analyze Tiering task is shown.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule > Analysis**.

**Related tasks:**

“Optimizing storage tiering” on page 269

To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

**Related reference:**

“Criteria for identifying the pools that require balancing” on page 278

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

## Tiering volumes by I/O rate

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To restrict the placement of volumes in tier 1 storage pools, you add minimum threshold values for the I/O rate.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is `svc_sales_records`, and the names of the DS8000 back-end storage systems are `ds81` and `ds82`. The pools on the storage virtualizer are assigned to tier 1, tier 2, and tier 3.

With the Analyze Tiering wizard, you want to analyze the placement of the volumes on tier 1. You set minimum threshold values for the I/O rate to re-tier the volumes that do not meet the minimum threshold criteria that you set for the I/O rate for each tier of storage.

On the Optimize the Placement of Volumes page of the Analyze Tiering wizard, you want to set the following minimum threshold for the I/O rate for each tier.

*Table 37. Minimum threshold values for tier 1, tier 2, and tier 3*

Tier	I/O rate minimum threshold level
1	5000
2	1000

Table 37. Minimum threshold values for tier 1, tier 2, and tier 3 (continued)

Tier	I/O rate minimum threshold level
3	500

The information in Table 38 table is provided to illustrate the re-tiering of the volumes that are currently on tier 1.

Table 38. Current I/O for the volumes in tier 1 pools

Pool	Current tier	Volumes	Current I/O rate
pool_ds81_1	1	volume_1	4000
pool_ds81_1	1	volume_2	1500
pool_ds81_1	1	volume_3	6000
pool_ds82_101	1	volume_11	7000
pool_ds82_101	1	volume_12	900
pool_ds82_101	1	volume_13	450

When you run the Analyze Tiering wizard, the I/O rate of the volumes on tier 1 is compared with the minimum threshold levels that you set for the storage tiers. The volumes in the tier 1 pools with values that are lower than the minimum threshold levels are selected as candidate volumes for re-tiering. The following criteria determine the destination pool for the candidate volumes:

- The minimum threshold levels that you set for tier 2 and tier 3 storage on the Optimize the Placement of Volumes page of the Analyze Tiering wizard
- The activity limits that you set for destination pools on the Optimize the Placement of Volumes page of the Analyze Tiering wizard.

In this scenario, the following recommendations are generated:

- Move volume\_1, and volume\_2 from tier 1 to a pool on tier 2.
- Move volume\_12 from tier 1 to a pool on tier 3.

Recommendations are not generated to re-tier volume\_3 and volume\_11 because the I/O rate of the volumes meets the minimum threshold levels for tier 1 storage. Recommendations are also not generated to re-tier volume\_13 because the I/O rate of the volume does not meet the minimum threshold levels that you set for tier 1, tier 2, or tier 3.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Double-click **svc\_sales\_records**.
3. To group the pools by tier, in the Internal Resources section, double-click **Pools**, and then click the **Tier** column heading.
4. Select all of the pools on tier 1, and then right-click and select **Analyze Tiering**.
5. Select all of the pools on tier 2 and tier 3 as target pools, and then click **Next**. To group the pools by tier, click the **Tier** column heading.
6. On the Optimize the Placement of Volumes page, select **Show advanced options**.
7. To specify the minimum thresholds for placing volumes on tier 1, complete the following steps:

- a. Set an activity limit for the tier. To ensure that the performance of the pools does not deteriorate when volumes are added, you set an activity limit for all of the pools on each tier.
- b. Select **I/O rate**, and type 5000.
8. To specify the minimum thresholds for placing volumes on tier 2, complete the following steps:
  - a. Set an activity limit for the tier.
  - b. Select **I/O rate**, and type 1000.
9. To specify the minimum thresholds for placing volumes on tier 3, complete the following steps:
  - a. Set an activity limit for the tier.
  - b. Select **I/O rate**, and type 500.
10. Click **Analyze**.

The Analyze Tiering task is shown.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule > Analysis**. Alternatively, you can modify the criteria that you set for analyzing the volumes on the Tasks page and run the analysis again.

**Related tasks:**

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

## Collocating volumes

In this scenario, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes in the Analyze Tiering wizard and in the Balance Pools wizard.

The volumes that you want to analyze are in storage pools on a Storwize V7000 Unified storage virtualizer with three DS8000 back-end storage systems. To ensure that volumes in the same storage pool that are assigned to the same server or hypervisor are kept together, you want to enforce the collocation of volumes. By enforcing the collocation of volumes, you prevent the placement of related volumes in destination pools that might be on multiple back-end storage systems.

**Multiple host connections to the same hypervisor or server:** If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

*Table 39. Information about the volumes before they are collocated*

Back-end storage system	Pool	Tier	Volume	Server
ds8000_1	pool_1	1	volume_1	server_1
ds8000_1	pool_1	1	volume_2	server_1
ds8000_1	pool_1	1	volume_3	server_2

Table 39. Information about the volumes before they are collocated (continued)

Back-end storage system	Pool	Tier	Volume	Server
ds8000_1	pool_1	1	volume_4	server_2
ds8000_2	pool_2	2	volume_5	server_3
ds8000_3	pool_3	3	volume_6	server_4

1. Log in to the Tivoli Storage Productivity Center web-based GUI.
2. In the menu bar, go to **Storage > Pools**.
3. Right-click the **pool\_1** storage pool, and then click **Analyze Tiering**.
4. Select **pool\_2** and **pool\_3** as the target storage pools.
5. On the Optimize the Placement of Volumes page, ensure that **Colocate volumes** is set to **Yes**.
6. Click **Show advanced options**, and choose the tiering criteria for analyzing the volumes.
7. Click **Analyze**.

If a volume requires re-tiering, a recommendation is generated to move all of the volumes that are assigned to the same server in the **pool\_1** storage pool to the same destination pool. If none of the target pools has sufficient space to accommodate all of the volumes, then recommendations to move the volumes are not generated.

**Related tasks:**

“Optimizing storage pools” on page 276

Use the Balance Pools wizard to analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

“Optimizing storage tiering” on page 269

To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

---

## Monitoring capacity usage at different levels of a business hierarchy

The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

The department model comprises these main elements:

- The department
- The subdepartments that the department contains
- The applications, or the application subcomponents, or both that the department uses

In the following scenario, you want to create a department model that monitors the capacity and space usage of the Books Sales department, and these subdepartments:

- Wholesale
- Retail

- Online

To monitor the capacity and space usage of the departments, you add the applications that the department and its subdepartment use to each of the subdepartments.

The department model that you create determines how you can view capacity usage and space information. To view the total capacity and space usage of a department regardless of the subdepartments that the department contains, you create a department and add the applications and subcomponents to the department.

If, however, you want to see the capacity and space usage of the department and its subdepartments, you create a department hierarchy by adding subdepartments to departments. You then add the applications and application subcomponents that each subdepartment uses to the subdepartments. You can then view capacity and space usage for the department and for the individual subdepartments.

You can create more complex models of departments by adding subdepartments to subdepartments. In the following illustration, the department model is extended to include two more layers of subdepartments under the Wholesale subdepartment.

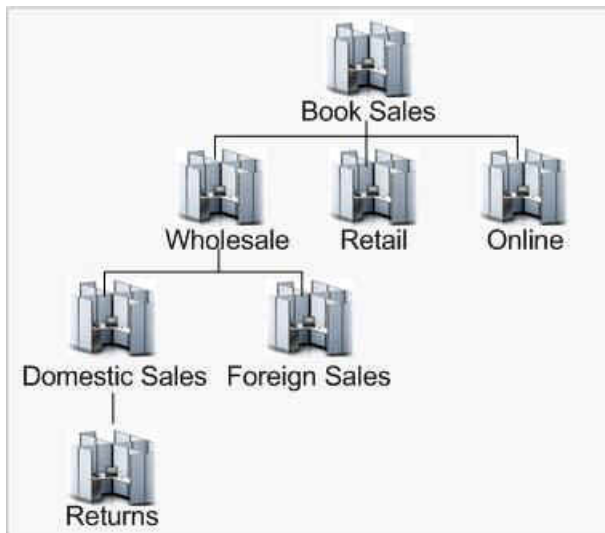


Figure 2. Extended department model hierarchy

To monitor the capacity and space usage of the departments in your business organization, complete the following tasks:

- Create the Book Sales department.
- Create the Wholesale, Retail, and Online departments to represent the subdepartments.
- Add the Wholesale, Retail, and Online as subdepartments to the Book Sales department.
- Create the Book Sales DB application and the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents that are used by the departments.

Different applications require different storage capabilities, and different levels of performance and uptime. The storage requirements of the departments are

constantly growing and they need to modify their behaviors in a way that justifies the cost of their actions concerning storage utilization. The data modeling of storage resources using the application and department concept enables you to plan and implement a chargeback system if necessary.

**Related concepts:**

“Departments” on page 233

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

**Related tasks:**

“Creating departments” on page 233

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in Tivoli Storage Productivity Center.

“Adding subdepartments” on page 234

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add subdepartments.

“Viewing information about departments” on page 235

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments (member departments) that are used to establish a business hierarchal structure.

“Creating applications” on page 225

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

“Adding subcomponents” on page 227

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

## Comparing storage usage in each department

In this scenario, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

To monitor the capacity and space usage of a department, you create the department and subdepartments, and you add the application subcomponents to the subdepartments.

You want to monitor the capacity and space usage of departments in an organization that sells books.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online




Figure 3. A department with three subdepartments hierarchy

Although you can add applications and application subcomponents when you create departments, it is easier and quicker to create the applications and application subcomponents beforehand.

You want to know the total amount of storage space that the Book Sales department uses, and you want to know the amount of space that each subdepartment uses. To know how much space the Book Sales and its subdepartments use, you must associate the department and its subdepartments with the applications that they use. For example, the Book Sales department and its subdepartments use a database application named Book Sales DB. The Book Sales DB application contains these subcomponents:

- Wholesale Transactions
- Retail Transactions
- Online Transactions

To create the department and subdepartments, you complete these tasks:

1. You create the Wholesale department and add the applications and application subcomponents that the subdepartment uses.
2. You create the Retail department and add the applications and application subcomponents that the subdepartment uses.
3. You create the Online department and add the applications and application subcomponents that the subdepartment uses.
4. You create the Book Sales department. You do not have to add applications or subcomponents to the Book Sales department because the Book Sales department inherits the applications and application subcomponents that are added to the subdepartments.
5. You add the Wholesale, Retail, and Online departments as subdepartments to the Book Sales department.

When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

You want to target what department and what storage resources might be responsible for a change or an upward trend in significant storage usage. By correctly classifying the department or departments you can properly plan to archive or switch the storage within the organization.

**Related concepts:**

“Departments” on page 233

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

**Related tasks:**

“Creating departments” on page 233

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in Tivoli Storage Productivity Center.

“Adding subdepartments” on page 234

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add subdepartments.

“Viewing information about departments” on page 235

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments (member departments) that are used to establish a business hierarchal structure.

“Creating applications” on page 225

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

“Adding subcomponents” on page 227

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

---

## Using applications and subcomponents to monitor capacity and space usage

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchal manner.

You can create a simple application model, which consists of an application and use a filter to associate the storage resources that the application uses with the application. Alternatively, you can create a complex application model that comprises an application and application subcomponents and then add filters to associate the storage resources that each application subcomponent uses with each of the subcomponents.

Besides creating applications to monitor the performance, capacity, and space usage of the applications in your business organization, you also create applications to add them to departments. When you add applications to departments, you can monitor the capacity and space usage of the departments and you can monitor the performance of the storage resources that are associated with the applications that each department uses.

To monitor the capacity of an application, you create an application model that is based on the following elements:

- The application



- The subcomponents that the application contains
- The applications, application subcomponents, or both that are added to the department

In the following scenario, you want to create an application model that monitors the capacity, space usage, and performance of the Book Sales DB application. The application model that you want to create comprises these application subcomponents:

- Wholesale Transactions
- Retail Transactions
- Online Transactions

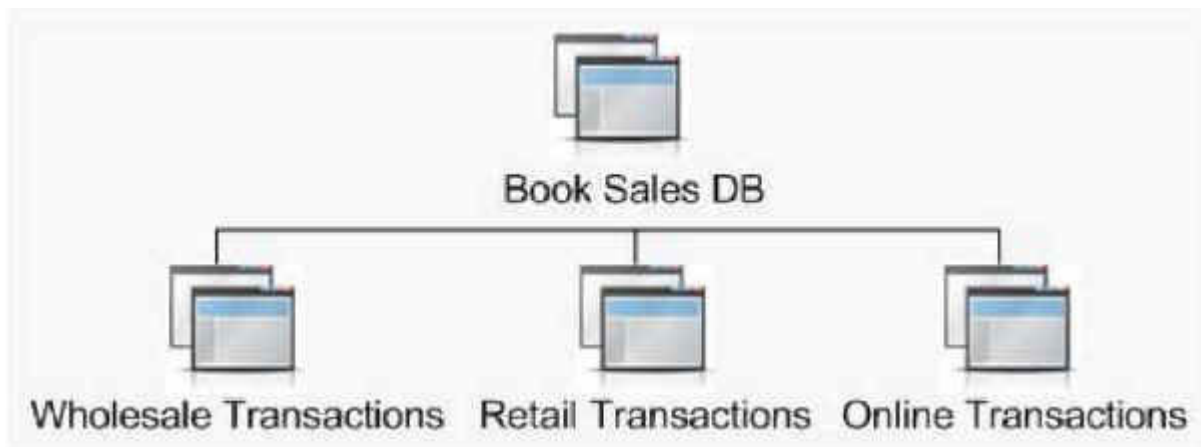


Figure 4. An application with a three subcomponent hierarchy

To understand how capacity is trending for storage in an application and subcomponents, complete these tasks:

1. Create the Book Sales DB application.
2. Create the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents.
3. Add the Wholesale Transactions, Retail Transactions, and Online Transactions as subcomponents to the Book Sales DB application.
4. Create the Book Sales department and the Wholesale, Retail, and Online subdepartments to be associated with the applications.

**Related concepts:**

“Applications” on page 225

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

**Related tasks:**

“Creating applications” on page 225

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

“Adding subcomponents” on page 227

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

“Viewing information about applications” on page 230

You can view information about the application that you created such as related resources, filters that are used to populate the application with resources, any subcomponents and capacity trending and performance information.

## Viewing storage capacity and usage trends

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

You want to understand how capacity is trending for storage that is used by the Book Sales DB application and the Wholesale Transaction, Retail Transactions and Online Transactions subcomponents and the association the application has with the Book Sales department and subdepartments.

The storage resources that are used by the application and that are used by each subcomponent are associated with the application and its subcomponents. For example, the storage resources are assigned as follows:

1. Create a database application called Book Sales DB and assign resources that use the volumes on the IBM® Storwize® V7000 with names that begin with **bksales** for storing sales transactions.
2. Create the subcomponent applications named Wholesale Transaction, Retail Transactions and Online Transactions.
3. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales\_ws** to the Wholesale Transactions subcomponent, which is used for storing wholesale sales transactions.
4. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales\_rt** to the Retail Transactions subcomponent, which is used for storing retail sales transactions.
5. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales\_ol** to the Online Transactions subcomponent, which is used for storing online sales transactions.

**Note:** On the Create Filter page, you can specify a name pattern to determine which volumes to include. For volumes, you can specify name patterns to determine from which servers, storage systems, or pools the volumes are selected. You can then click **Preview** to view the volumes that are selected for inclusion in your application.

To monitor space usage for the Book Sales department and its subdepartments, you associate the Book Sales DB application and the Wholesale, Retail, and Online Transactions subcomponents of the application as follows:

- You associate the Book Sales DB application with the Book Sales department.
- You associate the Wholesale subdepartment with the Wholesale Transactions subcomponent.
- You associate the Retail subdepartment with the Retail Transactions subcomponent.

- You associate the Online subdepartment with the Online Transactions subcomponent.

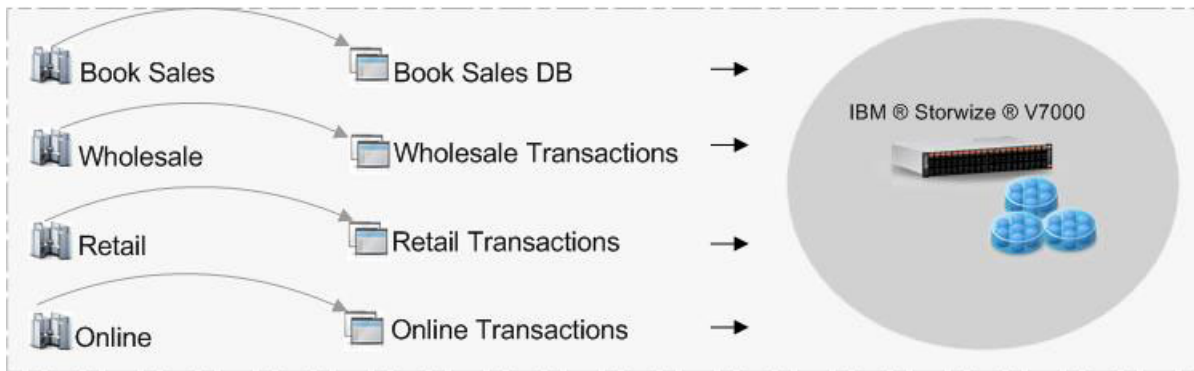


Figure 5. The application, subcomponent, assigned resources and department association

You can view the capacity information that is collected about the application on the Applications page. You can also view capacity information for each application and subcomponent on the details page. When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

**Related concepts:**

“Applications” on page 225

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

**Related tasks:**

“Creating applications” on page 225

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

“Adding subcomponents” on page 227

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

“Adding applications to departments” on page 227

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

“Viewing information about applications” on page 230

You can view information about the application that you created such as related resources, filters that are used to populate the application with resources, any subcomponents and capacity trending and performance information.

---

## Tracking performance issues on a storage system

You can be automatically notified of performance issues on storage system by defining alerts. Alerts can be triggered when the configuration of a storage system changes, or when it is not performing as expected. In this scenario, you define a performance threshold to be notified when the total I/O rate of a storage system falls outside an expected range.

To monitor the performance of a storage system, you must collect performance data and define alerts. To collect performance data, schedule a performance monitor. Performance monitors collect data about monitored resources and at each sample interval the data is compared to the performance thresholds that you specify.

In this scenario, you are defining a performance threshold alert to monitor a SAN Volume Controller named **SVC-01-IBM**. The metric for the threshold is total I/O rate, which measures the average number of I/O operations per second for both sequential and nonsequential read and write operations. After data is collected and the threshold is checked, you are viewing the threshold violation and exporting the performance information to a CSV file.

Complete the following tasks for this scenario:

- Scheduling data collection for a storage system.
- Defining a performance threshold and determine how you are notified when a violation is detected.
- Viewing information about the violations of a performance threshold.

### Scheduling data collection



Before you can check for performance issues on a storage system, you must set up data collection. Typically, data collection is automatically scheduled when resources are added to Tivoli Storage Productivity Center. To ensure that data collection was run, complete the following steps:

1. Log in to the web-based GUI.
2. In the menu bar, go to **Storage > Storage Systems**.
3. On the Storage Systems page, locate the row for the storage system.
4. Check the values in the **Probe Status** and **Performance Monitor Status** columns to ensure that data collection was successful.
5. If data collection was not run, complete the following steps:
  - a. To run a probe, right-click the storage system and select **Data Collection > Start Probe**.
  - b. To run a performance monitor, right-click the storage system and select **Data Collection > Start Performance Monitor**.

### Defining a performance threshold

Specify performance boundaries for the Total I/O Rate threshold and determine how you are notified when those boundaries are violated.


1. Log in to the web-based GUI.
2. In the menu bar, go to **Storage > Storage Systems**.
3. Right-click the storage system for which you want to define a performance threshold and click **Edit Alert Definitions**.
4. Click **Performance** from the category list.

5. Click the switch icon  that is displayed next to **Total I/O Rate Threshold** to enable it.
6. To define the boundary values for the threshold, click the performance chart icon  that is displayed next to it.
7. Enter boundary values for triggering warning and critical performance alerts. In this scenario, the following values are used:
  - a. In the **Critical > Greater than** field, type **3000**.
  - b. In the **Warning > Greater than** field, type **2200**.

**Tip:** These threshold values are examples and are not recommended values. They are for illustration purposes only. Performance behavior can vary significantly between different resources and even different applications for different resources. It is recommended that you monitor the performance of a resource for several weeks and use the historical data to determine reasonable values for a threshold.

8. Optional: To dynamically evaluate performance on the chart, click **Top 10** or **Bottom 10** to change the internal resources of the storage system that are displayed, click a time period to change the granularity, or use the date and time fields to customize the history of data. In the chart, the yellow horizontal lines represent the boundary values that you enter in the **Warning** fields. The red horizontal lines represent the boundary values that you enter in the **Critical** fields.



**Tip:** You can use the information from the performance chart to help you decide the boundary values to set for the threshold.

9. Click **Done**. After you define the performance threshold, determining how you are alerted to violations of performance thresholds is important to helping you maintain and administer storage resources. In this scenario, you set up the alert definition so that a violation of the threshold is automatically written to the operating system log.
10. Click the envelope icon  that is displayed next to **Total I/O Rate Threshold**.
11. Select **Override notification settings** to define notification settings that are specific to this threshold. If you do not select this option, the global notification settings that are defined for the storage system are used.
12. Select **UNIX syslog** to write alert notifications to the operating system log. On a Windows system, you can write alert notifications to the Windows log. Writing alerts to the operating system log is a way to centralize priority messages for quick notification and viewing.
13. Click **Done** to accept the notification settings.
14. Click **Save**.

## Viewing threshold violations

In this scenario, the total I/O rate for an internal resource exceeded 3000 for SVC-01-IBM and a threshold alert was generated. You want to open the alert for the threshold violation and view the performance of the resource that caused the violation.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. In the list of storage systems, double click the storage system (SVC-01-IBM) that generated an alert.

3. To view the alerts that were triggered for the storage system, click the **Alerts** tab.
4. In the list of alerts, double click the **Total I/O rate threshold** violation that you want to view. A chart shows the historical performance of SVC-01-IBM. The chart includes lines that represent the critical boundary and warning boundaries that were defined for the performance threshold. The default time range of the chart spans 2 hours before and 2 hours after the violation occurred.  
In this case, the total I/O rate for several volumes on SVC-01-IBM contributed to the threshold violation. A list of the volumes that were estimated to be the top contributors to the workload are shown below the chart.
5. Optional: To access the full performance view for SVC-01-IBM, click the open performance view icon . In the full performance view, you can select different metrics for the storage system, change the time range of the chart, view internal resources, and other tasks.
6. Optional: To export the performance information to a CSV file, click the export icon  on the chart.

**Tip:** You can use the preceding steps, with some minor changes, to create threshold alerts for certain types of monitored switches.

**Related tasks:**

“Creating performance monitors” on page 152

Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in the web-based GUI when you add resources for monitoring or at a later time.

“Defining a performance threshold alert for a resource” on page 164

Performance alerts are triggered by violations of performance thresholds on a storage system or switch. To specify the performance thresholds, you must edit the performance alert definition.

“Viewing performance information” on page 206

View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

**Related reference:**

“How performance information is displayed” on page 210

The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

---

## Limiting a set of storage to provision in the vSphere Web Client extension

In this scenario, a storage administrator limits the storage that a vSphere user can provision.

You plan to dedicate two storage pools to be used with vSphere. You want to allow a vSphere user to provision storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center. You want to define the two storage pools as the resources that the user can use.

1. Create a capacity pool that the vSphere user can provision from. Place the two storage pools that you designated as provisioning targets for vSphere in the

- capacity pool. To add storage pools to the capacity pool, follow the instructions in “Adding resources to capacity pools” on page 252.
2. Create a service class and specify that the service class can be used to provision storage only from the capacity pool. To create the service class, follow the instructions in “Creating service classes” on page 246.
  3. Create a user ID for the vSphere user and assign the External Application role to that ID. To assign the role to the user, follow the instructions that are provided in the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_t\\_associate\\_user\\_group\\_to\\_roles.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_t_associate_user_group_to_roles.html).
  4. Associate the user with the service class. To associate the user with the service class, follow the instructions in “Modifying service classes” on page 248.

When you complete these tasks, a vSphere user can provision storage by using the vSphere Web Client extension, but they are limited to the service class that was created. Because the service class is associated with the capacity pool that contains your designated vSphere storage pools, you limited where the vSphere user can provision data.

**Related concepts:**

“Cloud configuration” on page 70

To take advantage of the simplified provisioning and optimization capabilities that are available by using the web-based GUI, some configuration is required.

**Related tasks:**

“Provisioning block storage” on page 267

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

“Provisioning file storage” on page 268

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

## Viewing NPIV connections between server ports and switch ports in a fabric

In this scenario, a storage administrator views the N-Port ID Virtualization (NPIV) connections between the server ports and the switch port in the fabric. With NPIV, multiple node ports can be logically connected to one switch port.

You want to view connectivity information for a switch port that has NPIV connections.

1. In the menu bar in the web-based GUI, go to **Network > Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.

**Tip:** Alternatively, in the menu bar, go to **Network > Fabrics**. Double-click the fabric that you want to view, or right-click the fabric and click **View Details**. In the **Internal Resources** section, click **Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.

2. In the **Internal Resources** section, click **Ports**. You can view details about the ports, including the number of NPIV connections.

3. In the **Connected NPIV Ports** column for the port that you want to view, click the number of ports. You can view the following information about the NPIV connections:
  - In the **Connected NPIV Ports** tab, you can view details about the NPIV connections, such as the port type and the name of the connected resource. The port type can be N\_Port (end node port) or NP\_Port (proxy node port).
  - In the **Connectivity** tab, you can view details about the physically connected end-node port.

**Tip:** If Tivoli Storage Productivity Center does not differentiate between the physical and logical connections for the switch port, the details about the NPIV connections are displayed in the **Connectivity** tab. The **Connected NPIV Ports** tab is not displayed.

---

## Replication scenarios

These scenarios illustrate how to perform specific replication tasks using IBM Tivoli Storage Productivity Center for Replication.

The provided scenarios set out situations that can be solved using IBM Tivoli Storage Productivity Center. When using these scenarios you should remember the following:

- The scenarios are intended to be step-by-step instructions that you follow to complete a specific task. Each scenario contains multiple steps that must be followed in the order given.
- Specific values such as user names, passwords, and IP Address are given in the scenarios only for illustration purposes. You must replace them with values appropriate for your environment.
- The scenarios do not provide in-depth conceptual information about the tasks.

## Performing planned maintenance scenarios

Scenarios for a planned outage consist of the removal of a unit from service to perform scheduled work on specific components for a predetermined duration.

### Planned outage of local (H1), moving production to site 2 (H2) and optional return

In this scenario, your goal is to shut down the H1 box for maintenance, or for a planned power outage. As a result, you must move production to the H2 site, with the option of switching back to the H1 site after it is restored.

Follow these steps to carry out this scenario.

#### 1. Start H1->H2->H3

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. The session will reach a Prepared state so there is a consistent point on Sites 2 and 3.

**Note:** Since the outage is a planned outage, stop production I/O to H1.

#### 2. Suspend

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

#### 3. Recover H2



This command makes H2 target-available. When the session state is Target Available, the production I/O can then be moved to write to the H2 volumes.

4. **Start H2->H3**

This command starts a Global Mirror session during the outage of H1. Production I/O can continue on H2. After the session reaches the Prepared state, a consistent copy of the data resides on Site 3, and can be used in the event of a disaster.

5. **Start H2->H1->H3**

Issue this command when H1 is available and ready to be brought back into the configuration. This command starts a Metro Global Mirror session with H1 and H2 roles reversed from the original configuration. Production I/O can continue on H2. After the session reaches the Prepared state, both Site 1 and Site 3 will have a consistent copy of the data.

**Note:** The following steps are the steps you must follow after you decide to switch production I/O to write on Site 1 again. Stop production I/O on H2 before continuing to the next step.

6. **Suspend**

This command suspends the Metro Mirror relationship between H2 and H1. When this step is complete, the session state should be suspended and recoverable.

7. **Recover H1**

This command makes H1 target-available. When the session state is target-available at H1, run production I/O to H1.

8. **Start H1->H2->H3**

This command returns the session to its original configuration, copying data from Site 1, to Site 2, to Site 3.

## **Planned outage on site 2 (H2) and return to three-site configuration**

In this scenario, there is a planned outage expected on site 2. Disaster recovery capability is still needed so the copy will be set up between site 1 and site 3. When site 2 is available again, you can restart the session and use the original session configuration.

Follow these steps to carry out this scenario:

1. **Start H1->H2->H3**

Run this command to begin Metro Global Mirror on your system when the production I/O is running on H1. After the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

**Note:** Production I/O does not stop running to site 1 to perform the following steps.

2. **Start H1->H3**

This command starts a Global Mirror session between site 1 and site 3, during the outage of H2. Production I/O can continue on H1. After the session has reached a Prepared state, a consistent copy of the data will reside on Site 3.

**Note:**

- If running DS8000 devices on H1 with the correct level of microcode, and enough time has passed since the initial start of the session, the creation of the H1 to H3 pairs will *not* require a full copy.

- Data on Site 2 is no longer consistent at this point.
- After Site 2 is brought back up and is ready to hold a consistent copy of the data, go to the next step to return to the full three-site configuration. Production I/O never has to stop writing to H1 in this scenario.

### 3. Start H1->H2->H3

After the session reaches a Prepared state, site 2 and 3 will have a consistent copy of the data again.

## Planned outage of local (H1) with production move to remote (H3)

In this scenario, there is a planned outage for Site 1 (H1). Site 2, however, does not have the resources to run production, so production will be moved to Site 3 during the outage at Site 1. When the outage is complete, production can then be moved back to Site 1 and resume the full three-site configuration.

Complete the following steps to carry out this scenario.

### 1. Start H1->H2->H3

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

**Note:** Because this is a planned outage at Site 1, stop production I/O to Site 1 (H1) before continuing to the next step.

### 2. Suspend

This command creates the planned outage by suspending the Metro Mirror relationship between H1->H2. When this step is complete, the session state should be suspended and recoverable.

### 3. Recover H3

This command makes H3 Target Available. When the session state is Target Available at H3, production I/O can be moved to H3.

**Note:** When Site 1 is ready to resume production, continue with the following steps. Production I/O can continue to be written to H3 through the following step.

### 4. Start H3->H1->H2

Issue this command when H1 and H2 are available and ready to be brought back into the configuration. Changes from production at the remote site (H3) flow from the local site (H1) to the intermediate site (H2). The H3-H1 and H1-H2 pairs are Global Copy. The session will remain in a preparing state.

**Note:** While the session is Preparing, you can decrease the amount of time that production I/O has to be stopped by waiting for progress on the H3-H1 role pair to be at or near 100%. The next step will drive H3-H1 to sync, which depending on the percentage of the copy that is complete, can take a significant amount of time. When you are ready to switch back to Site 1, stop production I/O on H3 before continuing to the next step.

### 5. Suspend

This command suspends Global Copy relationship between H3->H1 and H1->H2. When this step is complete, the session state should be suspended and recoverable.

### 6. Recover H1

This command makes H1 Target Available. When the session state is Target Available at H1, production can be switched back to H1.

7. **Start H1->H2->H3**

This command will resume copy from Site 1 back to sites 2 and 3. Once the session reaches a Prepared state, there will be a consistent copy on Site 2 and Site 3.

### **Planned outage of remote (H3) and returning**

In this scenario, the remote site H3 has a planned outage, or consistency groups from the Global Mirror session need to be temporarily paused.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. After the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

**Note:** Production I/O to H1 can continue during the following steps.

2. **Suspend H2->H3**

Issue this command to suspend only the Global Mirror portion of the session. The Metro Mirror continues between H1 and H2. This causes the session state to be put into a SuspendedH2H3 state.

3. **Start H1->H2->H3**

This command restores the original configuration by restarting the Global Mirror portion of the session.

## **Recovering from a disaster scenarios**

Scenarios for disaster recovery (or unplanned outages) describe the steps to recover from the loss of a complete site and restart operations after a disaster or disruption.

### **Unplanned outage of local (H1), production move to intermediate (H2)**

In this scenario, there was an unplanned outage of the local site (H1), and production must be moved to Site 2 (H2) until Site 1 can be brought back up.

Complete the following steps to carry out this scenario.

**Note:** To ensure consistency in the event of a rolling disaster where the IBM Tivoli Storage Productivity Center for Replication server is lost on Site 1 before freezing all the copies in the session, enable the Metro Mirror heartbeat function.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

2. **Unplanned outage of H1**

This event causes an unplanned outage of H1, which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the state was Preparing when the outage occurred, the session might be recoverable at Site 2 but will not be recoverable at Site 3.

3. **RecoverH2**

This command makes H2 Target Available. When the session state is Target Available at H2, production can be switched to H2.

4. **Start H2->H3**

This command starts a Global Mirror session from Site 2 to Site 3 during the outage of H1. Once the session reaches a Prepared state, consistent data should reside on Site 3.

**Note:** Do not perform the following step until Site 1 is back up and ready to receive a copy of the data that has been written to Site 2 during the outage.

5. **Start H2->H1->H3**

This command starts an MGM session with the H1 and H2 roles reversed from the original configuration. Production I/O can continue on H2. Wait for the session to go to the Prepared state, at which point there will be a consistent copy of the data on Site 1 and Site 3.

**Note:** The following steps can be used to return to the original configuration so that production I/O can be switched back to Site 1. Stop production I/O on H2 before continuing to the next step.

6. **Suspend**

This command suspends the Metro Mirror relationship between H2 and H1. When this step is complete, the session state should be suspended and recoverable.

7. **RecoverH1**

This command makes H1 Target Available. When the session state is Target Available at H1, production I/O can be switched back to H1.

8. **Start H1->H2->H3**

This command restores the original configuration. Once the session reaches a Prepared state, a consistent copy of the data will reside on Site 2 and Site 3.

## Unplanned outage of Site 2 (H2)

In this scenario, an unplanned outage of Site 2 (H2) occurs. Disaster Recovery capability is still required, so a Global Mirror session must be set up between Site 1 (H1) and Site 3 (H3) while Site 2 is down.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data

**Note:** Production I/O to H1 can continue during the following steps.

2. **Unplanned outage of H2**

This event causes an unplanned outage of H2, which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the session state was Preparing, the session might be recoverable at Site 2 but would not be recoverable at Site 3.

3. **Start H1->H3**

This command starts a Global Mirror session between Site 1 and Site 3 during the outage of H2. Production I/O can continue on H1.

**Note:** If you are running DS8000 devices on H1 with the proper level of microcode, and enough time has passed since the initial start of the session, the creation of the H1 to H3 pairs will NOT require a full copy.

Once the session has reached a Prepared state, a consistent copy of the data will reside on Site 3.

**Note:** Data on Site 2 is no longer consistent at this point.

**Note:** After Site 2 has been successfully brought back up, perform the next step to return to the three-site configuration.

#### 4. Start H1->H2->H3

Issue this command when H2 is ready to be brought back into the configuration. This command restores the original configuration. Production I/O can continue on H1. Once the session reaches a Prepared state, a consistent copy will reside on Site 2 and Site 3.

### Unplanned outage of remote (H3)

In this scenario, an unplanned outage occurs on Site 3 (H3).

Complete the follow these steps to carry out this scenario.

#### 1. Start H1->H2->H3

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

A failure of the remote site (H3) causes the session state to go from Prepared to SuspendedH2H3, with an error message that the connection to the H3 box has been lost.

**Note:** When the remote site (H3) is repaired and ready to return to the configuration, the session will remain SuspendedH2H3 and recoverable until step 2 is completed.

#### 2. Start H1->H2->H3

Issue this command only when the remote site (H3) is repaired and ready to return to the configuration. With this command, your session returns to the original configuration by starting the Global Mirror portion back up.

### Unplanned outage of local (H1) with production move to remote (H3)

In this scenario, there is an unplanned outage for Site 1 (H1). Site 2, however, does not have the resources to run production, so production will be moved to Site 3 during the outage at Site 1. When the outage is complete, production can then be moved back to Site 1 and resume the full three-site configuration.

Complete the following steps to carry out this scenario.

**Note:** To ensure consistency in the event of a rolling disaster where the IBM Tivoli Storage Productivity Center for Replication server might be lost on Site 1 before freezing all the copies in the session, enable the Metro Mirror heartbeat function.

#### 1. Start H1->H2->H3

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. When the session is in a Prepared state, both sites 2 and 3 have a consistent copy of the data.

#### 2. Unplanned outage of H1 and H2

This event causes an unplanned outage of H1 and H2 which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the state was Preparing, the session might be recoverable at Site 2, but will not be recoverable at Site 3.

**3. Recover H3**

This command makes H3 Target Available. The session will be severe/Target Available. When the session state is Target Available at H3, production I/O can be moved to H3.

**Note:** Once Site 1 is brought back up, perform the following step to begin the process of returning production to Site 1.

**4. Start H3->H1->H2**

Issue this command when H1 and H2 are available and ready to be brought back into the configuration. Changes from production at the remote site (H3) flow from the local site (H1) to the intermediate site (H2). The H3-H1 and H1-H2 pairs are Global Copy. The session will in the preparing state.

**Note:** While the session is Preparing, you can decrease the amount of time that production I/O has to be stopped by waiting for the progress on the H3-H1 role pair to be at or near 100%. The next step will drive H3-H1 to sync, which depending on the percentage of the copy that is complete, can take a significant amount of time. When you are ready to switch back to Site 1, stop production I/O on H3 before continuing to next step.

**5. Suspend**

This command suspends the Global Copy relationship between H3>H1 and H1>H2. When this step is complete, the session state should be suspended and recoverable.

**6. Recover H1**

This command makes H1 Target Available. When the session state is Target Available at H1, run production I/O to H1.

**7. Start H1->H2->H3**

This command restores the original three-site configuration. Once the session reaches a Prepared state, a consistent copy of the data will reside on Site 2 and Site 3.

## Practicing disaster recovery scenarios

This section provides scenarios for how to practice disaster recovery using a regular Metro Global Mirror session. You can follow these scenarios to practice disaster recovery while continuing to be disaster recovery-capable using only the Metro Mirror portion of the session. The Metro Global Mirror with Practice session is also available, which allows for practicing disaster recovery while maintaining full three-site disaster recovery capability.

### Practice a recovery to H2 with production that is continuing on H1 in a normal configuration

In this scenario, your goal is to practice a recovery to H2 with production that is continuing on H1 in a normal configuration.

Follow these steps to carry out this scenario.

**1. Start H1->H2->H3**

Issue this command to begin a Metro Global Mirror session on your system when the production I/O is running on H1. Wait for the session to go to the prepared state.

You can continue to run production I/O to H1 during the following steps.

**2. Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

**Note:** Stop running production I/O to H2 prior to issuing the Suspend command.

**3. Recover H2**

This command makes H2 target-available. When the session is target-available at H2, run practice I/O to H2.

**4. Start H1->H3**

Issue this command to ensure that any I/O that is written to site 1 is copied to site 3.

You can also use the Start H1->H3 command to practice for a disaster at the H2 site while still maintaining disaster recovery capability. When the H1->H3 role pair is in Global Mirror mode, you can recover the session at site 3 if a disaster occurs.

When you complete testing on H2, you can issue the Start H1->H2->H3 command to return to a three-site configuration without completing a full copy.

**5. Start H1->H2->H3**

Issue this command to restart the Metro Global Mirror session in the H1-H2-H3 configuration. Any changes from production at the local site (H1) flow from the remote site (H3) to the intermediate site (H2), which overwrites the test data on H2.

Table 40. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

Steps	Role-pair states					
	H1-H2	H2-H3	H2-J3	H3-J3	H1-H3	H1-J3
1	Prepared	Preparing	Prepared	Target available	Non-participating	Non-participating
2	Suspended					
3						
4		Non-participating	Non-participating		Participating	Prepared
5	Prepared	Preparing	Prepared		Non-participating	Non-participating

**Practice recovery to H1 with production that is continuing on H2**

In this scenario, your goal is to practice a recovery to H1 with production that is continuing on H2.

Follow these steps to carry out this scenario.

**1. Start H1->H2->H3**

Issue this command to begin a Metro Global Mirror session on your system when the production I/O is running on H1. Wait for the session to go to the prepared state.

**2. Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

**Note:** Stop running production I/O to H1 prior to issuing the Suspend command.

**3. Recover H2**

This command makes H2 target-available.

**Note:** Start running production I/O to H2.

**4. Start H2->H1->H3**

This command starts a Metro Global Mirror session with H1 and H2 roles reversed from the original configuration.

You can continue to run production I/O to H2 for the next two steps.

**5. Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session should be suspended and recoverable.

**6. Recover H1**

This command makes H1 target-available. When the session state is target-available at H1, run practice I/O to H1.

**7. Start H2->H3**

Issue this command to ensure that any I/O that is written to site 2 is copied to site 3.

You can also use the Start H2->H3 command to practice for a disaster at the H1 site while still maintaining disaster recovery capability. When the H2->H3 role pair is in Global Mirror mode, you can recover the session at site 3 if a disaster occurs.

When you complete testing on H1, you can issue the Start H2->H1->H3 command to return to a three-site configuration without completing a full copy.

**8. Start H2->H1->H3**

Issue this command to restart the Metro Global Mirror session in the H2-H1-H3 configuration. Any changes from production at the local site (H2) flow from the remote site (H3) to the intermediate site (H1), which overwrites the test data on H1.

Table 41. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

Steps	Role-pair states					
	H1-H2	H2-H3	H2-J3	H3-J3	H1-H3	H1-J3
1	Prepared	Preparing	Prepared	Target available	Non-participating	Non-participating
2	Suspended					
3						
4	Prepared	Non-participating	Non-participating		Preparing	Prepared
5	Suspended					
6						
7		Preparing	Prepared		Non-participating	Non-participating
8	Prepared	Non-participating	Non-participating		Preparing	Prepared

**Practice at H3 for Metro Global Mirror**

This topic describes how to practice at the remote site (H3) while running a Metro Global Mirror session.



To practice on the remote site, you first need to suspend the back-end. For example, if running in an H1->H2->H3 configuration, issue a **Suspend H2H3** command to pause the Global Mirror portion of the session while leaving the Metro Mirror portion of the session active. When the session is in a suspended state (that is, H2H3 is suspended while in an H1->H2->H3 configuration), you have the option to issue a **Recover H3** command.

The **Recover H3** command acts differently than other recover commands. In this case, the command acts only to make the remote site (H3) target-available and does not switch the application site to H3. The application site remains the original host site and the session acts as it does in a normal Metro Mirror Failover/Failback session.

While the session is in this mode, it remains capable of disaster recovery at Site 2 and will freeze all pairs to maintain consistency during an error. You can also suspend the session by command and recover to Site 2 if necessary. However, if recovering to Site 2 in this mode, it will not be possible to restart a full three-site solution. To get three sites back up and running, you need to copy back to H1 using the **Start H2->H1** command and then suspend and recover to H1. Once H1 is the application host again, you can issue the **Start H1->H2->H3** command to restart three-site support. If this is done, all data written to H3 during the practice is overwritten with the data that is on H1.

### **Overview of scenarios for practicing at H3 for Metro Global Mirror:**

This topic provides an overview of the scenarios for practicing at H3 for Metro Global Mirror.

**Note:** For simplicity, the following scenarios are geared toward the configuration where the session is running H1->H2->H3. The same options apply when in the configuration is running H2->H1->H3.

There are two methods to set up a practice mode at the remote site, H3.

In the first method, you issue the following commands:

1. **Start H1->H2->H3**
2. **SuspendH2H3**
3. **RecoverH3**

When the session state is Target Available at H3, you can practice your disaster recovery steps on H3 while continuing to run the real production I/O to H1.

In the second method of setting up a practice mode, issue the following commands:

1. **Start H1->H2->H3**
2. **Suspend**
3. **RecoverH3**
4. **Start H1->H2**

When the session state is Target Available at H3, run practice I/O to H3. Continue running production I/O to H1.

The difference between the second method and the first method is that a **Suspend** command is issued instead of a **SuspendH2H3** command. When the session is in suspended mode (as opposed to suspended H2H3 mode), the session does not

automatically go into practice mode. The application site switches to H3 and you must issue a **Start H1->H2** to enable the practice mode.

**Practice scenario 1: While practicing, disaster occurs at Site 1, recover to Site 2:**

This topic describes what to do in the event that a disaster occurs on Site 1 while practicing at Site 3.

If a disaster occurs at Site 1, the session immediately suspends itself. You then must complete the following steps:

1. Issue **RecoverH2**.
2. Move the production I/O to Site 2.
3. Issue **Start H2->H1** to make the session capable of disaster recovery again once Site 1 has been brought back up.

**Practice scenario 2: While practicing, planned outage at Site 1, recover to Site 2:**

This topic describes what to do if there is a planned outage on Site 1 while practicing at Site 3.

While practicing, if you want a planned outage with a switch to Site 2, complete the following steps:

1. Issue **Suspend**.
2. Issue **RecoverH2**.
3. Move the production I/O to Site 2.
4. Issue **Start H2->H1** to make the session capable of disaster recovery again once Site 1 has been brought back up.

**Practice scenario 3: while practicing with production at site 2, move production back to site 1:**

This topic describes an H3 practice scenario, with production running at site 2, followed by a move of production to site 1.

While practicing, if you moved the production to site 2, you need to move it back to site 1 before you can start three-site support. To move production back to site 1, do the following steps (assume the session is running H2->H1):

1. Issue **Suspend**.
2. Issue **Recover H1**.
3. Move the production I/O to site 1.
4. Issue **Start H1->H2** to make the session capable of disaster recovery again and move the production I/O to H3.
5. Issue **Start H1->H2->H3** when H3 is ready to receive a copy of the data at site 1. When this command completes, all writes that were done to H3 are overwritten by the data on site 1.

The following table shows the changing states of the role pairs after each step.

Table 42. Role-pair states at the completion of each step (participating and nonparticipating role pairs)


Steps	Role-pair states					
	H1-H2	H2-H3	H2-J3	H3-J3	H1-H3	H1-J3
1	Suspended	Nonparticipating				
2	Target Available					
4	Prepared	target-available			nonparticipating	
5		Preparing	Prepared	Target Available		



---

## Chapter 5. Managing resources

Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

Tivoli Storage Productivity Center provides a web-based GUI and stand-alone GUI for monitoring the condition, capacity, performance, and relationships of storage resources. To view information about managing resources in the stand-alone GUI, go to the product documentation at  [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_managing.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_managing.html).

Tivoli Storage Productivity Center also includes Tivoli Storage Productivity Center for Replication. Tivoli Storage Productivity Center for Replication manages copy services in storage environments. *Copy services* are features that are used by storage systems such as System Storage DS8000 to configure, manage, and monitor data-copy functions. Copy services include IBM FlashCopy, Metro Mirror, Global Mirror, and Metro Global Mirror.

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage a vSphere virtual storage environment. You can complete provisioning tasks and view information about storage resources in the vSphere Web Client extension.

### **Related concepts:**

“Starting Tivoli Storage Productivity Center” on page 42

Tivoli Storage Productivity Center provides two graphical user interface (GUI) applications for managing and monitoring the resources in a storage environment: a web-based GUI and a stand-alone GUI. You can start these GUIs on the Tivoli Storage Productivity Center server or on a remote system.

“Provisioning storage with the vSphere Web Client extension” on page 266

You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage and monitor your virtual environment and storage devices.

### **Related reference:**

“Available functions in the interfaces” on page 34

Tivoli Storage Productivity Center provides the web-based GUI and stand-alone GUI for managing storage resources. In V5.2, V5.2.1, and V5.2.2, many of the functions for managing storage were moved to the web-based GUI. Determine which GUI to use based on the tasks that you must complete.

---

## Adding resources

You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

The type of resource that you want to add determines which Tivoli Storage Productivity Center GUI to use:

### **Web-based GUI**

The web-based GUI guides you through the steps for discovering the following resources in your environment, adding the resources as data sources, and scheduling data collection:

- Storage systems


- Servers
- Hypervisors
- Fabrics
- Switches

#### **Stand-alone GUI**

You can add the following resources in the stand-alone GUI:

- NAS filers
- Tivoli Storage Productivity Center servers as subordinate servers

#### **Tips:**

- For a complete list of the resources and their versions that you can add, see  [Tivoli Storage Productivity Center Interoperability Matrix](#).
- You cannot add resources that are already being monitored by Tivoli Storage Productivity Center. A resource is considered to be monitored if it is included in a data collection job such as a probe, performance monitor, or scan.
- Some types of storage systems and switches are managed by CIM agents. Tivoli Storage Productivity Center communicates with the CIM agent to discover those resources. During the discovery process, you can click **Close** to exit the wizard. When you click **Close**, the discovery process continues in the background, but any discovered resources are not configured for data collection. Before you can manage and view information about the discovered resources, you must schedule data collection.

#### **Related concepts:**

“Alerting” on page 159

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines resources for the conditions and performance thresholds that you specify and automatically notifies you when changes or violations are detected.

#### **Related reference:**

“Triggering conditions for alerts” on page 1036

Define alerts so that IBM Tivoli Storage Productivity Center automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

“Performance thresholds” on page 1100

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.











## **Resources that you can monitor in the web-based GUI**

You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

## Fabrics

A *fabric* is a network of hubs, switches, adapter endpoints, and connecting cables that support a communication protocol between devices. The following table shows the internal resources and related resources that you can monitor for fabrics.

Table 43. Fabric resources that you can monitor in the web-based GUI











Top-level resource icon	Internal resources	Related resources
	<ul style="list-style-type: none"> <li> Switches</li> <li> Inter-switch connections</li> <li> Blades</li> <li> Ports</li> <li> Zone sets</li> </ul>	<ul style="list-style-type: none"> <li> Hypervisors</li> <li> Servers</li> <li> Storage systems</li> <li> Discovered ports</li> </ul>

## Hypervisors

A *hypervisor* is a layer of software or a physical device that manages multiple, virtual instances of operating systems on the same hardware. Hypervisors that you can monitor in the web-based GUI include VMWare vSphere data sources such as ESX, ESXi, and vCenter servers. If a hypervisor is part of a cluster, information about that cluster is also available.

The following table shows the internal resources and related resources that you can monitor for hypervisors.












Table 44. Hypervisor resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	<ul style="list-style-type: none"> <li> Virtual machines</li> <li> Controllers</li> <li> Disks</li> <li> Data stores</li> <li> VMDKs</li> </ul>	<ul style="list-style-type: none"> <li> Storage systems</li> <li> Fabrics</li> <li> Switches</li> <li> Cluster</li> </ul>

## Servers

A *server* is a computer or host that is connected to a network and provides an access point to that network. If a server is part of a cluster, information about that cluster is also available. The following table shows the internal resources and related resources that you can monitor for servers.

Table 45. Server resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	<ul style="list-style-type: none"> <li> Controllers</li> <li> Disks</li> <li> Volume groups</li> <li> File systems and logical volumes</li> <li> Shares and exports</li> <li> Monitored directories</li> </ul>	<ul style="list-style-type: none"> <li> Storage systems</li> <li> Fabrics</li> <li> Switches</li> <li> VMDKs</li> </ul>

## Storage systems































A *storage system* is a device, such as a Redundant Array of Independent Disks (RAID) controller that creates and manages other storage devices. A storage virtualizer is a storage system that virtualizes storage space from internal storage or from another storage system.

Storage systems can be configured for block storage, file storage, or a combination of block and file storage. The method that a storage system uses for managing data determines the internal resources that are monitored.

The following table shows the internal resources and related resources that you can monitor for storage systems.















Table 46. Storage system resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	<p><b>Block storage resources</b></p> <ul style="list-style-type: none"> <li> Controllers</li> <li> Disks</li> <li> External disks</li> <li> Host connections</li> <li> I/O groups</li> <li> Managed disks</li> <li> Modules</li> <li> Nodes</li> <li> Pools</li> <li> Ports</li> <li> RAID arrays</li> <li> Volumes</li> </ul> <p><b>File storage resources</b></p> <ul style="list-style-type: none"> <li> Controllers</li> <li> Clusters</li> <li> Disks</li> <li> Shares</li> <li> Filesets</li> <li> File systems</li> <li> Network shared disks</li> <li> Nodes</li> <li> Pools</li> <li> Snapshots</li> </ul>	<ul style="list-style-type: none"> <li> Back-end storage systems</li> <li> Fabrics</li> <li> Hypervisors</li> <li> Servers</li> <li> Switches</li> <li> Virtualizer storage systems</li> <li> GPFS clusters</li> </ul>

## Switches

A *switch* is a network infrastructure component to which multiple nodes attach. In Tivoli Storage Productivity Center, you can monitor Fibre Channel (FC) switches. The following table shows the internal resources and related resources that you can monitor for switches.

Table 47. Switch resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	<ul style="list-style-type: none"> <li> Inter-switch connections</li> <li> Blades</li> <li> Logical switches</li> <li> Parent switches</li> <li> Ports</li> </ul>	<ul style="list-style-type: none"> <li> Fabrics</li> <li> Hypervisors</li> <li> Servers</li> <li> Storage systems</li> <li> Switches</li> <li> Discovered ports</li> </ul>

**Related tasks:**

“Adding ESX and ESXi hypervisors” on page 136

Add ESX and ESXi hypervisors for monitoring by Tivoli Storage Productivity Center.

“Adding storage systems”

To monitor a storage system, configure a connection to the storage system and schedule the collection of data.

“Adding fabrics and switches” on page 133

Add fabrics and switches for monitoring.

“Adding servers” on page 138

Add servers for monitoring. You can add a server by deploying a Storage Resource agent, or you can add an agentless server.

## Adding storage systems


To monitor a storage system, configure a connection to the storage system and schedule the collection of data.

To add a storage system, you must have Administrator privileges.

Use the web-based GUI to add the following types of storage systems for monitoring:

- DS8000
- XIV system
- DS6000
- ESS
- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V5000
- Storwize V7000
- Storwize V7000 Unified
- DS4000
- DS5000
- IBM SONAS
- GPFS clusters and GSS systems
- IBM FlashSystem V840 Enterprise Performance Solution

- All others (managed by CIM agents)

For a complete list of the versions of storage systems and CIM agents that you can add, see the  Tivoli Storage Productivity Center interoperability matrix and go to the Storage section.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Click **Add Storage System**.
3. Click the icon for the type of storage system that you want to add.

**Tip:** Click **Storwize Family** to add the following storage systems:

- IBM® Storwize® V3500
  - Storwize V3700
  - Storwize V5000
  - Storwize V7000
4. Complete the connection information for the storage system and schedule the collection of data.

The storage system is added for monitoring by Tivoli Storage Productivity Center. A probe is automatically run to collect status and asset information about the monitored resource.

Storage systems can be configured for block storage, file storage, or a combination of block and file storage. The method that a storage system uses for managing data determines the internal resources that are monitored.

When the collection of data is complete, you can view status information and capacity data about the storage system on the Storage Systems page.

## Adding fabrics and switches


Add fabrics and switches for monitoring.


For a switch to successfully receive and respond to queries from Tivoli Storage Productivity Center, the following basic requirements must be met:

- Tivoli Storage Productivity Center uses SNMPv1 to probe switches and fabrics, and uses SNMPv2 to collect performance data. Tivoli Storage Productivity Center does not support SNMPv3. Switches that Tivoli Storage Productivity Center probes must use the SNMPv1 protocol. Switches that Tivoli Storage Productivity Center collects performance data from must use the SNMPv2 protocol. Some switches are configured to use SNMPv3 by default.
- The Fibre Alliance FC Management MIB and Fibre Channel FE MIB must be enabled on the switch.
- The community string that is configured in Tivoli Storage Productivity Center must match one of the community strings that are configured on the switch with read access. Cisco switches must additionally have a community string match for write access. The default community strings in Tivoli Storage Productivity Center are "public" for read access and "private" for write access. Other community strings can be defined on the switches, but are not used.
- SNMP access control lists must include the Tivoli Storage Productivity Center system. Some lists automatically include all hosts while others exclude all by default.

In the web-based GUI, you can add switches and fabrics for monitoring at the same time. You can add the following types of switches:

- Brocade
- Cisco
- QLogic

For a complete list of the switches and their versions that you can add, check the switches and directors information in the  Tivoli Storage Productivity Center interoperability matrix.

**Help tips in the GUI:** To view descriptions of the information that you must enter for a fabric or switch, hover the mouse pointer over the related help icons  in the wizard.

1. In the menu bar in the web-based GUI, go to **Network > Switches**.
2. Click **Add Switch**.
3. Determine how to discover the switch and connected fabric.
  - Select **Use CIM agent as data source** to use a CIM agent as the data source. Tivoli Storage Productivity Center automatically discovers the switches that are managed by the agent. After this step, you can also select to use Simple Network Management Protocol (SNMP) agents for the switches.
  - Do not select **Use CIM agent as data source** to skip the use of a CIM agent as a data source. If you skip the use of a CIM agent as a data source, you are prompted to add switches and their fabrics through SNMP agents.

The type of switch and functions that you want to enable on that switch and connected fabric determine the agent to use for monitoring. For example, to monitor the performance of a switch, you must use either a CIM agent or an SNMP agent. To ensure full monitoring and functions for fabrics, you must also install Storage Resource agent on a server that is connected to the fabric.

For more information about agent types and the fabric and switch functions that they enable, see “Agent types for switch and fabric functions.”

4. Specify the display name and location of the switches and fabrics that are discovered. For switches that were discovered through a CIM agent, you can also select to monitor them through SNMP.
5. Schedule data collection for the switches and fabric. The type of resource determines the data collection jobs that you can schedule:
  - Fabrics: probes
  - Switches: probes, performance monitors

Probes collect status and asset information. Performance monitors collect metrics that measure performance.

6. Complete the wizard.

After a switch or fabric is added, data collection jobs gather information about them. When data collection is complete, you can view detailed information about the switch or fabric in the web-based GUI and stand-alone GUI.

### **Agent types for switch and fabric functions**

Depending on the functions that you want to enable, you can use a CIM agent, SNMP agent, or Storage Resource agent to monitor switches and fabrics.

The vendors of switches can help you determine which agents to use.

Table 48. Agent types for switch and fabric functions

Function > Switch	Brocade	Cisco <sup>1</sup>	QLogic
Switch Performance Monitoring	CIM agent or SNMP agent	CIM agent or SNMP agent	QLogic switches cannot be used for performance monitoring.
Switch and switch port information collected	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Topology connectivity information collected	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Zoning information collected <sup>2</sup>	Recommended: CIM agent	A Storage Resource agent is required in each VSAN <sup>3 4 5</sup>	Storage Resource agent
Tivoli Storage Productivity Center Alerts	Recommended: CIM agent Also supported: SNMP agent	SNMP agent	SNMP agent
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent	Storage Resource agent	Storage Resource agent
Switch "Sensors and Events"	SNMP agent	SNMP agent	SNMP agent
<p><b>Tips:</b></p> <ol style="list-style-type: none"> <li>1. The Storage Resource agent collects the information about a VSAN basis. Each VSAN is viewed as an individual SAN. The SNMP agent and CIM agent get the physical fabric information and can correlate the VSAN information to a physical infrastructure. The SNMP agent also collects some VSAN information.</li> <li>2. "Zone information collected" indicates that information about the active zone set and the inactive zoning library is collected during a fabric probe.</li> <li>3. For many Cisco proprietary zone members, the values are displayed as zeros.</li> <li>4. If a zoning configuration has Cisco proprietary zone members, then you cannot modify the zoning for that configuration through Tivoli Storage Productivity Center. You must use the element manager to remove proprietary zone members or do zoning.</li> <li>5. Orphan Cisco zones are not reported.</li> </ol>			


## Adding hypervisors

You can add hypervisors for monitoring by Tivoli Storage Productivity Center. Hypervisors include VMware vSphere data sources such as ESX and ESXi, and vCenter Server systems. When you add vSphere data sources, you can then collect data, generate reports, and manage storage that is related to those resources.

To add a hypervisor, you specify connection information for the ESX or ESXi data source. You can add multiple hypervisors by specifying connection information for a vCenter Server system.

## Adding ESX and ESXi hypervisors

Add ESX and ESXi hypervisors for monitoring by Tivoli Storage Productivity Center.

You can add an ESX or ESXi hypervisor by specifying connection information for that hypervisor. For a complete list of hypervisors that you can add, see the Agents, Servers and Browsers section in the  Tivoli Storage Productivity Center interoperability matrix.

1. In the menu bar in the web-based GUI, go to **Servers > Hypervisors..**
2. Click **Add Hypervisor**.
3. Select **ESX/ESXi** and enter connection information about the hypervisor.
4. Schedule a probe of the hypervisor.
5. Follow the instructions in the wizard to add the hypervisor.

After a hypervisor is added, a probe collects status and asset information about that hypervisor. You can view detailed information about the hypervisor in the web-based GUI and in the stand-alone GUI.

### Related tasks:

“Adding vCenter Server systems” on page 137

For Tivoli Storage Productivity Center to monitor multiple hypervisors, you must first add a vCenter Server system.

### Checking permissions to browse data stores:

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

When you add a VMware data source in Tivoli Storage Productivity Center, the user name that you specify must have permission to browse through the data stores on VMware. Tivoli Storage Productivity Center must browse through the data stores to collect information from the hypervisors. However, the "Read Only" role as defined by VMware does not allow Tivoli Storage Productivity Center to browse the data stores. You can use the "Virtual Machine Power User" role if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

To verify that a VMware user is assigned the correct role and privileges to monitor VMware data sources, follow these steps:

1. Ensure that the user role has the required VMware datastore permissions by completing the following steps:
  - a. Connect the vSphere Web Client to the VMware data source. The data source can be an ESX server, a vCenter Server Appliance, or a vCenter Server.
  - b. From the Inventories view, select **Hosts and Clusters**.
  - c. Select a host, and click the **Related Objects** tab.
  - d. View the datastores by clicking the **Datastores** tab.
  - e. Right-click a datastore, and select **File Browser**. If you can view the **Files** tab for the datastore, your browse permission is working correctly.
2. Determine the role that is assigned to the user by logging in to the vSphere Web Client by using the administrator user ID. From the Administration view, select **Roles**. Verify the role name that is assigned to the user.


3. Determine the privileges that are assigned to the role by selecting the user's role and clicking **Privileges**. Expand the privilege groups to view the specific privileges.
4. Optional: If you must edit the privileges for the role, select the role and click the **Edit role action** icon. Select privilege groups or expand to select specific privileges.

For more information about VMware user roles, go to the VMware documentation center and search for *vSphere users and permissions*.

## Adding vCenter Server systems

For Tivoli Storage Productivity Center to monitor multiple hypervisors, you must first add a vCenter Server system.

When you add vCenter Server version 5.1 or later, you can deploy the vSphere Web Client extension for Tivoli Storage Productivity Center and register the latter as a VASA provider. You can then provision storage, view reports, and publish alerts in the vSphere Web Client about storage that is monitored by Tivoli Storage Productivity Center.

You can add multiple hypervisors by specifying connection information for a vCenter Server system. For a complete list of hypervisors and vCenter Servers that you can add, see the Agents, Servers and Browsers section in the  Tivoli Storage Productivity Center interoperability matrix.

1. In the menu bar in the web-based GUI, go to **Servers > Hypervisors**.
2. Click **Add Hypervisor**.
3. Select **VMware vCenter** and enter connection information about the vCenter Server system.
4. Optional: To deploy the vSphere Web Client extension for Tivoli Storage Productivity Center and register Tivoli Storage Productivity Center as a VASA provider, in the **Deploy vSphere extension** page, enter your credentials. For **VMware vCenter Server**, enter a user name and password for the vSphere administrator user. For **Tivoli Storage Productivity Center**, enter a user name and password for a Tivoli Storage Productivity Center user ID. You must have Administrator, Monitor, or External Application authority to deploy the extension. If you entered a vSphere administrator user name and password for the vCenter Server system, enter only the Tivoli Storage Productivity Center credentials.
5. Schedule a probe for the hypervisors that were discovered.
6. Follow the instructions in the wizard to add the vCenter Server system.

After a vCenter Server system is added, probes collect status and asset information about the hypervisors. You can view detailed information about the hypervisors in the web-based GUI and in the stand-alone GUI.

### Related tasks:

“Adding ESX and ESXi hypervisors” on page 136

Add ESX and ESXi hypervisors for monitoring by Tivoli Storage Productivity Center.

“Checking permissions to browse data stores” on page 136

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

## Adding servers

Add servers for monitoring. You can add a server by deploying a Storage Resource agent, or you can add an agentless server.

Choose which type of server to add and which method you use to add the server.

To determine which type of server to add, follow these guidelines:

*Table 49. Choosing the type of server to add*

Condition	Add an agentless server	Deploy a Storage Resource agent to a server
You want to view the monitored resources to which a server is connected.	X	X
You want to view the amount of storage that the server is assigned from other monitored resources.	X	X
You want to view the performance of the SAN-related resources that are connected to the server.	X	X
The restrictions within your organization prohibit the deployment of agents to a server.	X	
You want to view information about the resources that are internal to a server. Internal resources include controllers, disks, volume groups, file systems, logical volumes, shares, exports, and directories.		X
You want to collect information about the fabrics that are visible to the server.		X

You can use several methods to add one or more servers for monitoring. Use the following information to determine the method that best meets the needs of your environment:

### Manually

Add one server by manually entering configuration information.

### From a File List

Add one or more servers by importing configuration information from a comma-delimited file.

### From Discovered Servers

Add one or more servers that are members of a Windows domain. If you select this method, the Add Server wizard automatically searches for servers that are in the Windows domain where Tivoli Storage Productivity Center is installed. When you select to add one server from the list of discovered servers, you can configure more information about that server.



When you select to add multiple servers, the servers are automatically created without additional configuration information.

**Restriction:** This method is available only when you add an agentless server. You cannot use this method when you add a server by deploying a Storage Resource agent.

### **Example: Server-centric monitoring of SAN resources without an agent**

In Tivoli Storage Productivity Center, you can view the connectivity of all your top-level resources in a SAN environment. These resources include servers, hypervisors, fabrics, switches, and storage systems. In previous releases, to get a complete view of the connectivity between servers and other top-level resources, you were required to deploy Storage Resource agents on all of those servers. However, in some environments it was not always possible to deploy Storage Resource agents, and thus a complete view of all resource relationships was not available. In Tivoli Storage Productivity Center V5.2 and later, you can add servers without deploying Storage Resource agents and still view the overall connectivity and performance of SAN storage that is assigned to those servers.

When you add an agentless server for monitoring and specify its worldwide port names, Tivoli Storage Productivity Center automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, the following information can be determined:

- The amount of SAN storage that is assigned to the server.
- The storage systems that are mapping volumes to the server.
- The relationships between the server and other monitored resources.

If performance monitors collect data about the related storage systems, you can also view the performance of the internal resources that are directly associated with storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.

### **Adding servers without Storage Resource agents**

Use the web-based GUI to add agentless servers for monitoring. An agentless server is a server or virtual machine that is monitored by Tivoli Storage Productivity Center but does not have a Storage Resource agent deployed to it.

When you add an agentless server, enter its IP address or host name and other optional information such as its OS type, location, and ports. Tivoli Storage Productivity Center uses this information to correlate the agentless server with the storage systems, fabrics, and switches that are being monitored. For example, the ports that you enter for an agentless server are compared to the ports on host connections for monitored storage systems. If there is a match between the server and a host connection, you can view the connectivity between the server and storage system and the capacity of volumes that are assigned to the server.

Use any of the following methods to add agentless servers for monitoring:

- Add a server by manually entering its information.
- Add a server from a host connection on a monitored storage system.
- Add a server from a discovered port that is connected to a monitored switch.

- Add a server from a discovered port that is connected to a monitored fabric.

**Add an agentless server by manually entering its information.**

1. In the menu bar in the web-based GUI, go to **Servers > Servers**.
2. Click **Add Server**.
3. Ensure that the **Deploy an agent for full server monitoring** check box is cleared.
4. Select a method for adding a server.
5. Complete the wizard.

**Add a server from a host connection on a monitored storage system.**

Use this method when you want to configure a host connection to represent a physical server. When you select a host connection, you are specifying the internal resources (volumes, ports) that belong to that server from the perspective of the related storage system.

To add an agentless server that is based on a known host connection, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Right-click a storage system and select **View Details**.
3. Click **Host Connections** to view a list of host connections that are mapped to a storage system.
4. Right-click a host connection and select **Add Server**. This action is not available under the following conditions:
  - A Storage Resource agent or agentless server is already associated with the host connection.
  - The host type of the host connection is configured for SAN Volume Controller.
5. Enter information about the server and assign ports.
6. Complete the wizard.

**Add a server from a discovered port that is connected to a monitored switch.**

Discovered ports are ports on resources that are not monitored by Tivoli Storage Productivity Center but are visible to a monitored switch. To add an agentless server that is based on a discovered port, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Network > Switches**.
2. Right-click a switch and select **View Details**.
3. Click **Discovered Ports** to view a list of unmanaged ports that are visible to the switch.
4. Right-click an unknown port and select **Add Server**.
5. Enter information about the server and assign ports.
6. Complete the wizard.


After the wizard is complete, the port is removed from the list of discovered ports and added as a server to the list of related resources for the switch.

**Add a server from a discovered port that is connected to a monitored fabric.**

Discovered ports are ports on resources that are not monitored by Tivoli Storage Productivity Center but are connected to ports on a monitored fabric. To add an agentless server that is based on a discovered port, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Network > Fabrics**.
2. Right-click a switch and select **View Details**.
3. Click **Discovered Ports** to view a list of unknown ports that are connected to ports in the fabric.
4. Right-click an unknown port and select **Add Server**.
5. Enter information about the server or virtual machine and assign ports.
6. Complete the wizard.

After the wizard is complete, the port is removed from the list of discovered ports and added as a server to the list of related resources for the fabric.

**Help tips in the GUI:** To view descriptions of the information that you must enter for a server, hover the mouse pointer over the related help icons  in the wizard.

**Restrictions:**

- The host name or IP address of an agentless server can be up to 24 characters in length and must contain valid characters. You can enter the following characters when you enter a host name or IP address:
  - A through Z (uppercase characters)
  - a through z (lowercase characters)
  - 0 through 9 (numeric characters)
  - Special characters: - (minus) . (period)
- If the location of an agentless server exceeds 64 characters, it is truncated when the server is added. To avoid the truncation, ensure that the location value for an agentless server is equal to or less than 64 characters in length.

**File List:**

Add one or more agentless servers by importing configuration information from a comma-delimited file. Tivoli Storage Productivity Center reads the entries in the file and automatically adds the corresponding servers.

The input file for importing server information can contain entries for one or more servers. Each line in the file represents a server that you want to add. Information about each server must be organized in the following format:

```
host name or IP address,OS type,machine type,location,WWPN
```

where:

- host name or IP address is required for each server entry. An IP address can be in an IPv4 or IPv6 format. You must include a host name or IP address for each server that you want to add. The host name or IP address can be up to 24 characters in length and must contain valid characters. You can enter the following characters when you enter a host name or IP address:
  - A through Z (uppercase characters)
  - a through z (lowercase characters)
  - 0 through 9 (numeric characters)
  - Special characters: - (minus) . (period)
- OS type is optional and represents the operating system of the server. The OS type for a server must be one of the following values:
  - Windows

- Linux
- AIX
- Solaris
- HP-UX
- Machine type is optional and determines if the server is a physical device or a virtual machine. The machine type for a server must be one of the following values:
  - Physical
  - Virtual
- Location is optional and represents the physical location of the server. If the location of an agentless server exceeds 64 characters, it is truncated when the server is added. To avoid the truncation, ensure that the location value for an agentless server is equal to or less than 64 characters in length.
- WWPN is optional and represents the worldwide port name that is associated with the server. You can include one or more WWPNs for each server. The value for World Wide Port Name (WWPN) must be 16 hexadecimal characters in length. If the WWPN for a server entry is already associated with a server that is being monitored, that WWPN is not created for the server.

#### Example

```
host1,Windows,Virtual,San Jose,2100000056AE45B2,2100000056AE45B3 # Comment
host5,Linux,Physical,London,2100000056EC45A4,21000000565E4FB6
systemxyz.storage.widgets.com,Solaris,,Chicago,21000000567E5FB7
198.51.100.22,HPUX,Virtual,,21000000566E6FB8
2001:DB8:0:0:0:0:0:0,Windows,Physical,Tokyo,21000000569E7FB9
```

**Tip:** If you add many servers in the same input file, the adding process might require some time to complete. To run the process in the background, click **Close** on the Add Server window. If you remain on the Servers page, a message is displayed when the servers are added. If you leave the Servers page during processing, a message is not displayed. To confirm that the servers were added, you must go back to the Servers page and refresh the list.

#### Restrictions

Keep in mind the following restrictions when you import server information from an input file:

- Each line in the file represents one server.
- Enter a "#" at the beginning of a line to comment out a line. The server on that line is not added when the list is imported in the add server wizard.
- If there are syntax problems in the file, then none of the servers in the file are added.

#### Adding servers with Storage Resource agents

Use the web-based GUI to add servers by deploying Storage Resource agents. Deploy Storage Resource agents if you want to enable full server monitoring.

To add a server, you must have Administrator privileges.

For each installation of Tivoli Storage Productivity Center, you can deploy only one Storage Resource agent on each server. If you attempt to deploy additional Storage Resource agents on a server, the deployments fail.

1. In the menu bar in the web-based GUI, go to **Servers > Servers**.
2. Click **Add Server**.

3. Select **Deploy an agent for full server monitoring**.
4. Select a method for adding a server. You can choose one of the following methods:
  - Add a server by manually entering information about the server and the Storage Resource agent.
  - Add one or more servers by importing configuration information from a comma-delimited file.
5. On the Deploy Agent page, configure deployment information for the Storage Resource agents.
 

If you add multiple servers with different operating systems, separate configuration pages are displayed for agents that are deployed on Windows servers and agents that are deployed on UNIX servers.
6. On the Configure page, schedule the deployment of the Storage Resource agents.
 

If you are deploying agents on multiple servers, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the Tivoli Storage Productivity Center server.
7. Schedule the time and frequency that probes are run for the servers.
 

If you add multiple servers, a time span is calculated during which the servers are probed.
8. To add the servers, click **Finish**.

A probe is automatically run for a server after the agent is successfully deployed.

Use the Status column on the Servers page to monitor the status of the agent deployment.

**File List:**

Add one or more servers with Storage Resource agents by importing the configuration information from a comma-delimited file.

The web-based GUI guides you through the following steps for adding servers:

- Select the input file.
- Configure deployment information.
- Schedule the agent deployment and data collection for the servers.

The comma-delimited file that you use to import the configuration information for servers can contain entries for a single server or multiple servers. Each line in the file represents a server that you want to add. The information about each server must be organized in the following format:

```
host name or IP address,OS type,location,custom tag 1,custom tag 2,custom tag 3
```

where:

- host name or IP address is required for each server entry. An IP address can be in an IPv4 or IPv6 format. A host name or IP address can contain the following characters or symbols:
  - A - Z (uppercase characters)
  - a - z (lowercase characters)
  - 0 - 9 (numeric characters)

- Symbols: - . : \_
- OS type is required and represents the operating system of the server. The OS type for a server must be one of the following values:
  - Windows
  - Linux
  - AIX
  - Solaris
  - HP-UX
- location is optional and represents the physical location of the server. The location value can be up to 64 characters in length. If the length exceeds 64 characters, the location value is truncated when the server is added.
- custom tag 1, custom tag 2, and custom tag 3 are optional and represent any additional information that you want to provide about the server. The custom tag values can be up to 64 characters in length. If the length exceeds 64 characters, the custom tag value is truncated when the server is added.

**Tip:** The custom tags can be displayed on the Servers page or can be included as report columns when you generate reports for the server.

**Example**

```
host1,Windows,San Jose,Accounting department
host5,Linux,London,Finance department
198.51.100.22,HP-UX,,Computing department
2001:DB8:0:0:0:0:0:0,Windows,Tokyo
```

**Tips:**

- If the comma-delimited file contains entries for multiple servers, it might take some time to add the servers. To confirm that the servers are added, check the Status column on the Servers page.
- To comment out a line, enter a "#" at the beginning of the line. The server on that line is not added when the list is imported.

**Example:**

```
# host1,Windows,San Jose,Accounting department
```

- If there are syntax problems in the file, none of the servers in the file are added.

## Removing resources

Remove resources that you no longer want to monitor with Tivoli Storage Productivity Center.

The type of resource that you want to remove determines which Tivoli Storage Productivity Center GUI to use:

**Web-based GUI**

Use the web-based GUI to remove the following resources:

- Storage systems
- Servers
- Hypervisors
- Fabrics
- Switches

**Stand-alone GUI**

Use the stand-alone GUI to remove the following resources:

- NAS filers
  - Tivoli Storage Productivity Center servers as subordinate servers
1. In the menu bar, go to the type of resource that you want to remove. For example, if you want to remove a server, go to **Servers > Servers**.
  2. On the list page for the resource, right-click the resource and click **Remove**.
  3. Click **Remove** to confirm that you want to remove the resource.

When you remove a resource, it is no longer monitored by Tivoli Storage Productivity Center. All the data that is associated with the resource, including historical data, is removed from the database. All the scheduled tasks and data collection jobs for the resource are canceled.

## Removing storage systems

Remove storage systems that you no longer want to monitor with Tivoli Storage Productivity Center.

To remove a storage system, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**. Information about monitored storage systems is displayed.
2. On the **Block Storage** or **File Storage** tab, right-click a storage system and click **Remove**.
3. Click **Remove** to confirm that you want to remove the storage system.

---

## Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. Tivoli Storage Productivity Center provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

Data collection jobs include probes, performance monitors, scans, and pings.

### Tips:

- The type of data collection job that you want to schedule determines which Tivoli Storage Productivity Center GUI to use.
- In the web-based GUI, data collection jobs are automatically scheduled when resources are added for monitoring. In the stand-alone GUI, you must schedule data collection jobs after resources are added.

Table 50. Data collection jobs, resources, and the GUI where you define the jobs









Data collection job	Resources	GUI
<p><b>Probe</b> Use probes to collect status, asset, and capacity information about resources.</p>	<p>Storage systems Servers Hypervisors Fabrics Switches</p>	<p>Web-based GUI:</p> <ul style="list-style-type: none"> <li>• Schedule a probe when you add a resource for monitoring.</li> <li>• Modify the schedule for a probe at any time on the list and details pages for a resource. For example, to modify the schedule for a storage system, go to the Storage Systems page or the Storage System details page.</li> </ul> <p> Learn more...</p>
<p><b>Performance monitor</b> Use performance monitors to collect metrics that measure the performance of storage systems and switches.</p>	<p>Storage systems Switches</p>	<p>Web-based GUI:</p> <ul style="list-style-type: none"> <li>• Schedule a performance monitor when you add a resource for monitoring.</li> <li>• Modify the schedule for a performance monitor at any time on the Performance page, and on the list and details pages for the resource.</li> </ul> <p> Learn more...</p>
<p><b>Scan</b> Use scans to collect file usage and trending information about servers or hypervisors.</p>	<p>Servers Hypervisors</p>	<p>Stand-alone GUI:</p> <p>Schedule a scan after a resource is added for monitoring. You can define a scan in the <b>Data Manager &gt; Monitoring &gt; Scans</b> node in the navigation tree of the stand-alone GUI.</p> <p>  Learn more...</p>
<p><b>Ping</b> Use pings to collect information about the availability (uptime) of servers.</p>	<p>Servers</p>	<p>Stand-alone GUI:</p> <p>Schedule a scan after a resource is added for monitoring. You can define a ping in the <b>Data Manager &gt; Monitoring &gt; Pings</b> node in the navigation tree of the stand-alone GUI.</p> <p>  Learn more...</p>



Table 50. Data collection jobs, resources, and the GUI where you define the jobs (continued)

Data collection job	Resources	GUI
<p><b>TPC Server probe</b> Use TPC server probes to gather the information that was collected by multiple Tivoli Storage Productivity Center servers.</p>	Tivoli Storage Productivity Center servers	<p>Stand-alone GUI:</p> <p>You can define a TPC Server probe in the <b>IBM® Tivoli Storage Productivity Center &gt; Monitoring &gt; TPC Server Probes</b> node in the navigation tree of the stand-alone GUI.</p> <p>  Learn more...</p>

## Collecting asset and status data

Use probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource list and resource details pages, for example, the Storage Systems page, or the Storage Systems details page.

### Creating probes

Probes are data collection jobs that collect status and asset information about your storage resources. You can create probes in the web-based GUI when you add resources for monitoring or after you add the resources for monitoring.

You can create probes for one or more resources at the same time.

To create probes for resources in your environment, you can use the list and detail pages for the resources. For example, to create probes for storage systems, use the Storage Systems page.

**Tip:** When you create a probe for a switch, the probe gathers statistics about the fabric that the switch is a part of.

1. From the menu bar in the web-based GUI, go to the resource list page for a resource type. For example, to create a probe for a storage system, go to **Storage > Storage Systems**.
2. Locate the resources that you want to probe.
3. Select the resource rows and click **Actions > Data Collection > Schedule**.
4. On the Data Collection Schedule window, set the probe to **Enabled**.
5. Enter the time of the data collection, select the frequency, and then click **Save**. If you are scheduling probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the server. A message is displayed that confirms that probe schedules are created for the resources.

Probe schedules are created for the resources that are selected.

The probe jobs run according to the modified schedule. If you disabled the probes, Tivoli Storage Productivity Center stops collecting asset and status data for the monitored resources.

You select three storage systems, storage\_system\_a, storage\_system\_b, and storage\_system\_c. A probe schedule was defined for storage\_system\_c, but not for storage\_system\_a and storage\_system\_b. You set the time and frequency fields, and set the probe to **Enabled**. When you click **Save**, probe schedules are created for storage\_system\_a and storage\_system\_b. The probe schedule that was defined for storage\_system\_c remains unchanged.

## Verifying that a probe is running for a resource

Verify that a probe is running for a resource to ensure that Tivoli Storage Productivity Center is collecting asset and status information about the resource.

To verify that a probe is running for a resource, use the resource list page in the web-based GUI.

1. From the menu bar, go to a resource list page for a resource type. For example, to verify that a probe is running for a storage system, go to **Storage > Storage Systems**.
2. Locate the resource that you want to check.
3. View the Probe Status column. A value of **Successful** or **Running** indicates that the probe is collecting asset and status information about the resource. The **Last Successful Probe** column shows the most recent date and time when the probe successfully collected data.
4. Optional: If the Probe Status column shows a status of **Failed** or **Warning**, right-click the resource and click **Data Collection > Open Probe Logs** to view the messages in the log file.
5. Optional: If the Probe Status column shows a status of **Never probed**, right-click the resource and click **Data Collection > Schedule** to schedule a probe for the resource.

## Modifying probes

Probes are data collection jobs that collect status and asset information about monitored resources in your environment. To modify a probe in Tivoli Storage Productivity Center, use the resource list page for the monitored resource in the web-based GUI.

You can modify the schedule for multiple probes at the same time.

1. From the menu bar, go to the resource list page for a resource type. For example, to modify a probe schedule for a storage system, go to **Storage > Storage Systems**.
2. Locate the resources that you want to probe.
3. Select the resource rows and click **Actions > Data Collection > Schedule**.
4. On the Data Collection Schedule window, modify the schedule for the probes, and click **Save**.

If you are scheduling probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the Tivoli Storage Productivity Center server.

The following rules are used to determine the settings for the probe fields when you open the Data Collection Schedule window:

- If probes are enabled for all of the selected resources, the **Enabled** value is displayed. If probes are disabled for all of the selected resources, the **Disabled** value is displayed. Otherwise, the state field for the probes is blank.

- For the time and frequency fields, if the selected resources have the same value for the field, the value is displayed. If the resources have different values for the field, the field is blank.

The probe jobs run according to the modified schedule. If you disabled the probes, Tivoli Storage Productivity Center stops collecting asset and status data for the monitored resources.

The changes that you make to the probe schedule are applied to all the resources that you select. For example, you select two servers with different probe frequency values. Probes are disabled for both servers. When you open the Data Collection Schedule window, the probe status field is set to **Disabled** and the frequency field is blank. You set the status field to **Enabled** and the frequency field to Every 2 days. When you click **Save**, the probe is enabled with the new frequency value for the two servers that you selected

### Configuring alerts for probes

Configure the alert notifications that are generated if a probe fails to run.

When a probe is scheduled for a resource in the web-based GUI, an entry for that probe is created in the stand-alone GUI. The entry is displayed in the navigation tree, under **IBM Tivoli Storage Productivity Center > Monitoring > Probe Alert Configurations**. The name of the entry is *userID.Probe\_resourcename*, where *userID* represents the user ID who created the probe and *resourcename* is the name of the resource that is being probed. For example, "user1.Probe\_DS8000a" represents a node for a probe that was created by user1 for a storage system named DS8000a.

You can edit the alert configuration for a resource if you want to modify how you are notified when a probe fails and which actions are taken. By default, if a probe fails, an error is recorded in the alert log. If you want to specify more actions if a probe fails for a resource, such as receiving an email and automatically running a script, you can modify its alert configuration.

**Tip:** You can view the alert log on the **Home > Alerts** page in the web-based GUI.

To modify the alert configuration for a probe, complete the following steps:

1. Start the stand-alone GUI.
2. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Monitoring > Probe Alert Configurations**.
3. Click the node for the alert configuration that you want to modify. For example, if you want to modify the alert configuration for the probe of a storage system named "Storwize\_IBM02", click the node **user1.Probe\_Storwize\_IBM02**.
4. In the **Triggered actions** section, determine the actions that occur as a result of a probe that fails. For information about triggered actions, see "Triggering actions and notifications for alerts" on page 1064.
5. To save the alert configuration, click **File > Save**.

If a probe for the resource fails, Tivoli Storage Productivity Center takes the actions that are defined in the alert configuration for that resource.

### Starting probes

To immediately collect data about resources, use the **Start Probe** action in the web-based GUI.

If a probe job is not defined for the resource, use the **Schedule** action to create a probe.

You can start probes for multiple resources at the same time.

**Tip:** When you start a probe for a switch, the probe gathers statistics about the fabric that the switch is a part of.

1. From the menu bar, go to the resource list page for a resource type. For example, to start a probe for a storage system, go to **Storage > Storage Systems**.
2. Locate the resources that you want to probe.
3. Select the resource rows and click **Actions > Data Collection > Start Probe**. The **Start Probe** action is available for the resources that you select if at least one of the resources meets the following criteria:
  - A probe job is defined for the resource.
  - The probe job is enabled.
  - The probe job is not currently running.

The probe status changes to Running.

Probes are started for all the resources that meet the criteria. For example, you select two storage systems, Storage\_System\_a and Storage\_System\_b. Probes are enabled for both of the storage systems and a probe is running for Storage\_System\_a. The probe is started for Storage\_System\_b when you select **Start Probe**. Storage\_System\_a is not affected by the action because a probe is already running for the resource.

### Viewing probes for a specific resource

You can view information for specific resources on the resource details pages. For example, use the Storage Systems details page to view the date and time when data was last collected.

1. In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want to view the resource information for a storage system, go to **Storage > Storage Systems**.
2. Locate the resource that you want to view information for.
3. Right-click the resource row and select **View Details**.
4. In the General section of the resource details page, click **Data Collection**.

The Data Collection pane is displayed, which contains information about the data collection jobs that are defined for the resource.





### Viewing probe logs

A probe log file is created for each run of a probe on a storage resource. Probe logs include detailed information about the status, actions, and progress of a probe.

1. In the menu bar in the web-based GUI, go to a resource list page for a resource type. For example, to view the probe logs for a storage system, go to **Storage > Storage Systems**.
2. Locate the resource that you want to view.
3. Right-click the resource row, and select **Data Collection > Open Probe Logs**.

**Restriction:** If you select multiple rows, the **Open Probe Logs** action is not available.

The Logs page opens and shows log entries for the most recent run of the probe.

4. Optional: On the Logs page, to view a previous probe log, select a log from the **Select a log** list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only entries that have the following statuses:
  -  Only error entries
  -  Only warning entries
  -   Error and warning entries
6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

You can view the following information on the Logs page:

- The overall status of a probe. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the job run.
- The status for each probe action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The description of an action.

## Collecting performance data

Before you can monitor the performance of the resources in your environment, you must collect data about those resources. Tivoli Storage Productivity Center uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the Performance Monitors page to view and manage performance monitors and to track threshold violations.

You must schedule and run performance monitors before you can complete other tasks, such as optimizing storage tiering, balancing pools, and running performance reports.

Before you can view performance information for resources, you must complete the following tasks:

- Add the storage system or switch for monitoring by Tivoli Storage Productivity Center.
- Schedule a performance monitor to collect performance data about the resource.
- Define a threshold in an alert so that you are notified when the performance of a resource might represent a potential problem. For example, you can define an alert threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range. You can define alerts in the stand-alone GUI.

The performance information that is gathered by Tivoli Storage Productivity Center includes metrics that measure the performance of volumes, ports, and disks. You can view performance metrics in the web-based GUI and stand-alone GUI at the following locations:

- Resource list and detail pages for switches and storage systems. For example, you can view performance metrics that are related to switches on the Switches page in the web-based GUI.
- Performance reports in the web-based GUI and stand-alone GUI. For example, you can access reports from the web-based GUI that show data rates, I/O rates, response times, and many other performance metrics.

### Creating performance monitors

Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in the web-based GUI when you add resources for monitoring or at a later time.

To create performance monitors, you must have Administrator privileges.

To create performance monitors for resources in your environment, you can use the list and detail pages for the resources. For example, to create performance monitors for switches, use the Switches page. You can create performance monitors for multiple resources at the same time.

#### Restrictions:

- You can define only one performance monitor for a resource.
- You can create performance monitors for the following resources only:
  - Storage systems
  - Switches that are managed by a CIM agent

1. From the menu bar, go to a resource list page for a resource type. For example, to create performance monitors for storage systems, go to **Storage > Storage Systems**.
2. Locate the resources that you want to monitor.
3. Select the resource rows and click **Actions > Data Collection > Schedule**.
4. On the Data Collection Schedule window, set the performance monitors to **Enabled**. You can set a performance monitor to **Enabled** only if the probe is set to **Enabled**.
5. Select the interval for the performance monitors, and click **Save**. The interval represents the number of minutes over which samples of performance data are averaged. If you are creating performance monitors for multiple resources, the interval list displays the performance monitor intervals that are common to all resources. For example, you select two storage systems. `Storage_System_a` specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. `Storage_System_b` specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

#### Restrictions:

Note the following restrictions when you create performance monitors:

- a. If you specify a 1-minute interval for performance monitors, data collection samples might be inconsistent lengths. The inconsistent lengths occur because it might take longer than 1 minute to collect the data on some systems. The inconsistent lengths can cause the data in the historical line chart to be misleading. This problem can also occur for the 5-minute summary chart when a 1-minute interval was specified for data collection. Check the messages in the job log to determine the length of time that data collection takes for a resource. If the data collection takes longer than 1

minute, consider changing the interval to 5 minutes. You can use a 1-minute interval when you want to troubleshoot a resource.

- b. If you specify a 1-minute interval, the amount of data that is stored in the database repository increases significantly. Tivoli Storage Productivity Center stores only 7 days of sample data that is collected at 1-minute intervals.

Performance monitors are created for the resources that are selected, and that meet the criteria. For example, you select three storage systems, `Storage_System_a`, `Storage_System_b`, and `Storage_System_c`. You enable the performance monitors and set the interval field. When you click **Save**, performance monitors are created for `Storage_System_a` and `Storage_System_b`. `Storage_System_c` is not affected by the action because you cannot configure performance monitors for that storage system.

If a successful probe run is completed for a resource, the performance monitor runs according to the defined interval.

To check the progress of a performance monitor, you can complete the following actions:

- From the resource list page, right-click a resource row, and select **Data Collection > Open Performance Monitor Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.
- View the Performance Monitor Status column on the resource list page.

### **Verifying that a performance monitor is running for a resource**

Verify that a performance monitor is running for a resource to ensure that Tivoli Storage Productivity Center is collecting performance metrics about the resource. Use performance monitors to collect metrics about the performance of storage systems and switches.

To verify that a performance monitor is running for a resource, use the resource list page in the web-based GUI.

1. From the menu bar, go to a resource list page for a resource type. For example, to verify that a performance monitor is running for a storage system, go to **Storage > Storage Systems**.
2. Locate the resource that you want to check.
3. View the Performance Monitor Status column. A value of Completed or Running indicates that the performance monitor is collecting performance metrics about the resource. The Last Successful Monitor column shows the most recent date and time when the performance monitor successfully collected data.
4. Optional: If the Performance Monitor Status column shows a status of Running with problems, Completed with warnings, Failed, or Canceled, you can view the logs to troubleshoot the error or warning messages. To view the performance monitor logs, right-click the resource row and click **Data Collection > Open Performance Monitor Logs**.
5. Optional: If the Performance Monitor Status column shows a status of Disabled, to enable the performance monitor for the resource, right-click the resource row and click **Data Collection > Schedule**. On the Data Collection Schedule page, select **Enabled**.

**Tip:** You can set the performance monitor to **Enabled** only if the probe is set to **Enabled**. If a probe is set to **Disabled**, the performance monitor cannot run.

## Modifying performance monitors

To modify all of the performance monitors that are used by Tivoli Storage Productivity Center, use the Performance Monitors page in the web-based GUI. To modify performance monitors for specific resources, use the resource list and resource detail pages.

To modify performance monitors, you must have Administrator privileges.

You can modify the schedule for multiple performance monitors at the same time.

You can enable or disable performance monitors and modify the interval. The interval represents the number of minutes over which samples of performance data are averaged.

1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**. Detailed information about all performance monitors is shown in the **Performance Monitors** tab.
2. Locate the performance monitors that you want to modify.
3. Select the performance monitor rows and click **Actions > Schedule**.
4. On the Data Collection Schedule window, modify the schedule for the performance monitors.

If you select multiple performance monitors, the following rules are used to determine the field settings when you open the Data Collection Schedule window:

- If all of the performance monitors that you select are enabled, the **Enabled** value is displayed. If all of the performance monitors are disabled, the **Disabled** value is displayed. Otherwise, the state field for the performance monitors is blank.
- If the performance monitors have different interval values, the interval field is blank. If the performance monitors have the same interval value, the value is displayed.

Intervals that are common to all performance monitors that you select are displayed in the interval list. For example, you select performance monitors for two storage systems. `Storage_System_a` specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. `Storage_System_b` specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

5. Click **Save**.

The changes are applied to all the performance monitors that are selected. The performance monitors run according to the modified schedule. If you disabled a performance monitor, Tivoli Storage Productivity Center stops collecting performance metrics for the storage system or switch.

To check the progress of the performance monitors, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the most recent date and time when performance data was collected about a resource.



- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor job. Use this information to troubleshoot any errors that might occur when the job runs.

## Configuring alerts for performance monitors

Configure the alert notifications that are generated if a performance monitor fails to run.

When a performance monitor is scheduled for a switch or storage system in the web-based GUI, an entry for that performance monitor is created in the stand-alone GUI. The entry is displayed in the navigation tree. For a storage system, an entry is displayed under **Disk Manager > Monitoring > Subsystem Performance Alert Configurations**. For a switch, an entry is displayed under **Fabric Manager > Monitoring > Switch Performance Alert Configurations**.

The name of the entry is *userID.Performance\_Monitor\_resourcename*, where *userID* represents the user ID who created the performance monitor and *resourcename* is the name of the associated resource. For example, "user1.Performance\_Monitor\_brocade1" represents a node for a performance monitor that was created by user1 for a switch named brocade1.

You can edit the alert configuration for a resource if you want to modify how you are notified when a performance monitor fails and which actions are taken. By default, if a performance monitor fails, an error is recorded in the alert log. If you want to specify more actions if a performance monitor fails for a resource, such as receiving an email and automatically running a script, you can modify its alert configuration.

**Tip:** You can view the alert log on the **Home > Alerts** page in the web-based GUI.

To modify the alert configuration for a performance monitor, complete the following steps:

1. Start the stand-alone GUI.
2. Choose one of the following options:

Option	Description
<b>Storage systems</b>	In the navigation tree, expand <b>Disk Manager &gt; Monitoring &gt; Subsystem Performance Alert Configurations</b> .
<b>Switches</b>	In the navigation tree, expand <b>Fabric Manager &gt; Monitoring &gt; Switch Performance Alert Configurations</b> .

3. Click the node for the alert configuration that you want to modify. For example, if you want to modify the alert configuration for the performance monitor of a switch named "switch\_brocade02", click the node **user1.Performance\_Monitor\_switch\_brocade02**.
4. In the **Triggered actions** section, determine the actions that occur as a result of a performance monitor that fails. For information about triggered actions, see "Triggering actions and notifications for alerts" on page 1064.
5. To save the alert configuration, click **File > Save**.

If a performance monitor for the resource fails, Tivoli Storage Productivity Center takes the actions that are defined in the alert configuration for that resource.

## Starting and stopping performance monitors

To start or stop performance monitors for storage systems and switches, use the Performance Monitors page in the web-based GUI.

To start or stop performance monitors, you must have Administrator privileges.

You can start or stop performance monitors for multiple resources at the same time.

**Restriction:** A successful probe must be completed on the resource before a performance monitor can start. If a successful probe is not run, you cannot start the performance monitor for the resource.

1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**.
2. Click the **Performance Monitors** tab.
3. Locate the performance monitors that you want to start or stop.
4. Select the performance monitor rows and click **Actions > Start** or **Actions > Stop**.

Performance monitors are started or stopped for all the resources that meet the criteria. For example, you select two storage systems, Storage\_System\_a and Storage\_System\_b. Performance monitors are enabled for the storage systems and both resources are probed. A performance monitor is running for Storage\_System\_a but not for Storage\_System\_b. The performance monitor is started for Storage\_System\_b when you select **Start**. Storage\_System\_a is not affected by the action because a performance monitor is already running for the resource.

Starting or stopping a performance monitor has the following effects:

**Start** Starting a performance monitor starts the immediate collection of data from the monitored resource. The status of the performance monitor changes to Running.

**Stop** Stopping a performance monitor stops the collection of performance data from the resource. If a performance monitor is running, it stops. Performance data is not collected from the resource until you restart the performance monitor. The status of the performance monitor changes to Canceled.

To check the progress of the performance monitor, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the total number of threshold violations that were detected during the last sampled time of the performance monitor.
- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.

## Viewing performance monitors

To view all of the performance monitors that Tivoli Storage Productivity Center uses to collect metrics for measuring the performance of resources, use the Performance Monitors page. To view the performance monitors for specific resources, use resource list and resource details pages.

## Viewing all performance monitors:

To view all of the performance monitors that Tivoli Storage Productivity Center uses to collect metrics for measuring the performance of storage systems and switches, use the Performance Monitors page in the web-based GUI.

The Performance Monitors page displays one row for each storage system or switch that Tivoli Storage Productivity Center is monitoring. If a performance monitor is not scheduled for a resource, the status column is set to Disabled for the performance monitor row.

Before Tivoli Storage Productivity Center can collect performance metrics for a resource, the resource must fulfill the following conditions:

- A probe must be run for the resource.
  - For switches, the switch must be managed by a CIM agent.
1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**. On the **Performance Monitors** tab, the information about performance monitors is organized into columns. These columns include the status of performance monitors, the names of the monitored resources, and the most recent dates and times when data was successfully collected.
  2. Optional: View the status icons at the top of the page to view a summary of performance monitor statuses. This summary includes the number of performance monitors that have a Running, Running with problems, Failed, or Not Running status.

## Viewing performance monitors for specific resources:

To view the performance monitors for specific storage systems and switches, use the resource detail pages in the web-based GUI. For example, use the Storage Systems detail page to view performance monitor information for a specific storage system.

If performance monitor information is not displayed on the resource detail page, this indicates that the resource does not support the collection of performance metrics by Tivoli Storage Productivity Center.

1. In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want to view the performance monitor for a storage system, go to **Storage > Storage Systems**.
2. Right-click a resource in the list, and select **View Details**. The resource detail page is displayed.

In the **General** section, the number next to **Data Collection** represents the number of data collection jobs that are associated with the resource. For example, **Data Collection (2)** indicates that you can view a performance monitor and a probe for the resource.

The icon next to the **Data Collection** link represents the most critical status of the data collection jobs that are associated with the resource. For example, if a job failed, the following icon is shown:

**Data Collection (2)** 

3. In the **General** section, click **Data Collection** to view information about the performance monitor for the switch or storage system. On the **Data Collection** pane, you can view information such as the status of the most recent run of a performance monitor, the performance monitor interval, and the most recent data and time that data was successfully collected.

**Tip:** You can use the columns on the resource list pages to view information about the performance monitors that are defined for storage systems or switches. For example, use the Storage Systems page to view performance monitor information for storage systems.





### Viewing performance monitors logs:

Use performance monitor logs to view detailed information about the status, actions, and progress of a performance monitor. You can use this information to troubleshoot any errors that might occur when a performance monitor is running.

Performance monitor logs contain informational, warning, and error messages that are related to each action that is taken during the processing of a performance monitor.

1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**.
2. In the **Performance Monitors** tab, locate the performance monitor that you want to view.
3. Right-click the performance monitor row, and select **Open Logs**.

**Restriction:** If you select multiple rows, the **Open Logs** action is not available. The Logs page opens and shows log entries for the most recent performance monitor log. A new log file is created when you stop and restart a performance monitor.

4. Optional: On the Logs page, to view a previous performance monitor log, select a log from the **Select a log** list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only entries that have the following statuses:
  -  Only error entries
  -  Only warning entries
  -   Error and warning entries
6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a performance monitor. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the job run.
- The status for each performance monitor action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

**Tip:** You can use the navigation and search functions of the web browser to locate information on the page.

## Collecting information about shares on storage systems

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Add **saveNonRoot=1** to the TPCD.config file, to view the information that is collected about shares in the web-based GUI.

1. Stop the Data server.
2. Open the TPCD.config file.

### For Windows operating systems:

By default, the TPCD.config file is in the C:\Program Files\IBM\TPC\data\config\ directory.

### For Linux or UNIX operating systems:

By default, the TPCD.config file is in either the usr/IBM/TPC/data/config/ directory or the opt/IBM/TPC/data/config directory.

3. In the [server] section of the TPCD.config file, add the following line of code:  
saveNonRoot=1
4. Save the TPCD.config file.
5. Restart the Data server.

When you provision new shares, you can view information about the shares on the Provisioning task page in the web-based GUI. On the Shares page, you can view more information about shares, such as the servers that have access to the share.

---

## Alerting

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines resources for the conditions and performance thresholds that you specify and automatically notifies you when changes or violations are detected.

Conditions that generate alerts are detected during data collection and event processing. For some storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system, events are polled every minute from the resource. For the GPFS/GSS storage system, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.


The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some triggering conditions, like performance thresholds, require you to enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- A change occurred in the storage infrastructure
- A performance threshold was violated
- A data collection job did not complete

For a complete list of the triggering conditions for alerts, see “Triggering conditions for alerts” on page 1036.


## Defining alerts



To set up alert definitions so that you are notified when certain conditions are detected on monitored resources, you must complete the following tasks:

Tasks	GUI	 <b>Learn more</b>
Adding resources for monitoring by Tivoli Storage Productivity Center.	Web-based GUI	“Adding resources” on page 127
Specifying the conditions that trigger alerts.	Web-based GUI	“Defining alerts” on page 162
Configuring how you are notified when alert conditions are detected.	Web-based GUI	“Viewing and modifying alert notification settings” on page 174
Collecting data about resources.	Web-based GUI: schedule probes and performance monitors  Stand-alone GUI: schedule pings and scans	Some alerts are triggered when conditions are detected during data collection. To detect when the configuration of a resource changes, run probes and scans. To detect when performance thresholds are violated, run performance monitors.  “Collecting data” on page 145  <b>Tip:</b> In the web-based GUI, probes and performance monitors can be scheduled when you add resources for monitoring.

## Viewing and managing alerts and alert definitions

After alert definitions are set up and alerts are being triggered, you can complete the following tasks to manage those alert definitions and triggered alerts:

Tasks	GUI	 <b>Learn more</b>
Viewing, acknowledging, exporting, and removing alerts. For alerts that are triggered by performance threshold violations, you can view a chart of the historical performance of the resource that violated a threshold.	Web-based GUI	<ul style="list-style-type: none"> <li>• “Viewing all alerts” on page 181</li> <li>• “Removing alerts” on page 184</li> <li>• “Acknowledging alerts” on page 184</li> </ul>

Tasks	GUI	 <b>Learn more</b>
Disabling alert definitions.	Web-based GUI	You can disable an alert definition so that the associated condition no longer generates alerts when it is detected. For example, to disable an alert for a storage system, go to the <b>Storage &gt; Storage Systems</b> page, right-click a storage system, and select <b>Edit Alert Definitions</b> . Click the switch icon next to a definition to disable it:  . For more information, see “Disabling or enabling alert definitions” on page 178.

#### Related tasks:

“Adding resources” on page 127

You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

#### Related reference:

“Triggering conditions for alerts” on page 1036

Define alerts so that IBM Tivoli Storage Productivity Center automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

“Performance thresholds” on page 1100

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

## Alert severities

Tivoli Storage Productivity Center determines the severity of alert conditions that it detects on monitored resources. Use the severity level to help determine the priority in which you resolve alerts.

Each severity is represented by an icon in the web-based GUI.

Table 51. Possible severities of alerts







Icon	Status	Description
	Critical	An alert with a Critical severity represents a serious problem on a resource or on its internal resources. Resolve these problems as soon as possible.
	Critical - Acknowledged	An alert with a Critical severity was acknowledged. A Critical - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Warning	An alert with a Warning severity represents potential problems on a resource or on its internal resources. The condition that generated the alert is not critical.
	Warning - Acknowledged	An alert with a Warning severity was acknowledged. A Warning - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.

Table 51. Possible severities of alerts (continued)

Icon	Status	Description
	Informational	An alert with an Informational severity does not represent a problem, but is intended to provide information about actions related to a resource.
	Informational - Acknowledged	An alert with an Informational severity was acknowledged. An Informational - Acknowledged severity indicates that an alert was reviewed and can be ignored.

## Defining alerts

Define how and when you are alerted to conditions and violations of performance thresholds on monitored resources. You can edit the alert definitions for a single resource or for multiple resources at the same time.

### Defining a configuration change alert for a resource

You can define alerts that are triggered when configuration changes are made to resources.

- To select a resource, choose one of the following options from the menu bar:


Option	Description
From the Storage menu, click Storage Systems.	Choose this option to select a storage system.
From the Servers menu, click Servers.	Choose this option to select a server.
From the Servers menu, click Hypervisors.	Choose this option to select a hypervisor.
From the Network menu, click Switches.	Choose this option to select a switch.
From the Network menu, click Fabrics.	Choose this option to select a fabric.

- To define a configuration change alert for the top-level resource, choose one of the following options:

Option	Description
Click the resource in the list, and then click Actions > Edit Alert Definitions.	Use this option to define the alert for a resource from the resource list page.
Right-click the resource in the list and click Edit Alert Definitions.	Use this option to define the alert for a resource from the resource list page.
Right-click the resource and click View Details. The resource page is displayed. In the General section, click Alerts, and then click the Definitions tab. Click Edit.	Use this option to edit the alert definition from the resource page.



**Tip:** You can define an alert for any of the resources in the category list. For example, to define an alert for storage system volumes, click **Volumes**.

- To enable the alert, click the switch icon .

**Tip:** To disable the alert, click the switch icon . When you disable an alert, notifications are not generated.

- For an alert condition with threshold values, enter the boundary values.






Option	Description
Click the switch icon  .	Use this option to enable the alert and set the boundary values.
Click the boundary values icon  .	Use this option to edit the boundary values for an enabled alert. <b>Tip:</b> Hover the mouse over the boundary values icon to see the threshold values for an enabled alert.

- Specify the threshold values.
- Click **Done**.
- Click **Save**.

### Defining a configuration change alert for multiple resources

You can define alerts that are triggered when configuration changes are made to multiple resources.

When you select multiple resources, the switch icon for each alert condition indicates how the condition is applied to the resources:

Icon	Status of the alert definition
	The alert definition is disabled.
	The alert definition is enabled.
	The alert definition is enabled for some of the selected resources and disabled for the other selected resources.

- To select resources, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click <b>Storage Systems</b> .	Choose this option to select storage systems.
From the Servers menu, click <b>Servers</b> .	Choose this option to select servers.
From the Servers menu, click <b>Hypervisors</b> .	Choose this option to select hypervisors.
From the Network menu, click <b>Switches</b> .	Choose this option to select switches.
From the Network menu, click <b>Fabrics</b> .	Choose this option to select fabrics.


- To select multiple resources, press Ctrl and click each resource.
- Click **Actions > Edit Alert Definitions**.


**Tip:** You can define an alert for any of the resources in the category list. For example, to define an alert for storage system volumes, click **Volumes**.

- To enable the alert, click the switch icon .

**Tip:** To disable the alert, click the switch icon . When you disable an alert, notifications are not generated.

- For an alert condition with threshold values, enter the boundary values.

Option	Description
Click the switch icon  .	Use this option to enable the alert and set the boundary values.

Option	Description
Click the boundary values icon  .	Use this option to edit the boundary values for an enabled alert. <b>Tip:</b> Hover the mouse over the boundary values icon to see the threshold values for an enabled alert.

6. Specify the threshold values.
7. Click **Done**.
8. Click **Save**.

### Defining a performance threshold alert for a resource

Performance alerts are triggered by violations of performance thresholds on a storage system or switch. To specify the performance thresholds, you must edit the performance alert definition.

When you define a performance threshold, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, a threshold violation occurs. For example, you can define a performance threshold that alerts you when the overall back-end response times for managed disks on a SAN Volume Controller exceed a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

**Tip:** Thresholds are checked against all the internal resources of a specific type; you cannot apply different thresholds to internal resources of the same type. For example, the storage system has two ports, Port A and Port B. You cannot set different threshold values for each port on the storage system. You must set the same threshold values for all of the ports on the storage system.



1. To select a resource, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click <b>Storage Systems</b> .	Choose this option to select a storage system.
From the Network menu, click <b>Switches</b> .	Choose this option to select a switch.

2. To edit the definition for a specific alert for a resource, choose one of the following options:

Option	Description
Click the resource in the list, and then click <b>Actions &gt; Edit Alert Definitions</b> .	Use this option to define the alert for a resource from the resource list page.
Right-click the resource in the list and click <b>Edit Alert Definitions</b> .	Use this option to edit the alert for a resource from the resource list page.
Right-click the resource and click <b>View Details</b> . The resource page is displayed. In the General section, click <b>Alerts</b> , and then click the <b>Definitions</b> tab. Click <b>Edit</b> .	Use this option to edit the alert definition from the resource page.

3. From the list of categories, click **Performance**.
4. To define a performance alert, choose one of the following options:

Option	Description
Click the switch icon  .	Use this option to enable the alert and set the boundary values.
Click the performance chart icon  .	Use this option to edit the boundary values for an enabled alert. <b>Tip:</b> Hover the mouse over the performance chart icon to see the boundary values for an enabled alert.

5. Enter boundary values for triggering warning performance alerts. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than 500 ops/s or less than 200 ops/s, in the **Warning > Greater than** field, enter the value 500. In the **Warning > Less than** field, enter the value 200.
6. Enter boundary values for triggering critical performance alerts. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than 600 ops/s and less than 100 ops/s, in the **Critical > Greater than** field, enter the value 600. In the **Critical > Less than** field, enter the value 100.

**Tips:**

- Tivoli Storage Productivity Center provides recommended values for thresholds that do not vary much between environments. For example, the default critical and stress boundaries for the disk utilization threshold are 80% and 50%. However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine boundary values for a resource, collect performance data over a period of time to establish a base-line of the normal and expected performance behavior for that resource. After you determine a set of base-line values, define thresholds to trigger if the measured performance behavior falls outside the normally expected range.
  - For some metrics and performance thresholds, lower values indicate more stress and higher values indicate idle behavior. For example, a lower value in the critical stress boundary for the Cache Holding Time Threshold might indicate a performance problem.
  - If you do not want to monitor violations for certain boundaries, leave those boundaries blank.
7. To evaluate performance information for the resource, click **Top 10** or **Bottom 10** to view the chart for resources ordered by performance, click a time period, or use the date and time fields to customize the performance chart display. In the chart, the yellow horizontal lines represent the values that you enter in the **Warning** fields. The red horizontal lines represent the values that you enter in the **Critical** fields.

**Tip:** You can use the information from the performance chart to help you decide the boundary values to set for the alert.

8. To filter any boundary violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
9. Click **Done**.
10. Click **Save**.

To determine whether performance thresholds are violated, the performance metrics that are collected for resources are compared with the threshold values that you defined for the alerts. If a violation occurs, you can view information about that violation in the following locations in the web-based GUI:

*Table 52. Locations where you can view violation of performance thresholds*

Location in GUI	Description
Storage Systems page Switches page	To view all the violations of performance thresholds on storage systems, go to <b>Storage &gt; Storage Systems</b> and click <b>Threshold Violations</b> .  To view all the violations of performance thresholds on switches, go to <b>Storage &gt; Switches</b> and click <b>Threshold Violations</b> .
Details page for a specific storage systems Details page for a specific switch	To view all the violations of performance thresholds on a specific storage systems, go to <b>Storage &gt; Storage Systems</b> , double click the storage system that you want to view, and click <b>Alerts</b> on the details page.  To view all the violations of performance thresholds on a specific switches, go to <b>Storage &gt; Switches</b> , double click the switch that you want to view, and click <b>Alerts</b> on the details page.
Performance Monitors page	To view the violations of performance thresholds for all resources, go to <b>Home &gt; Performance Monitors</b> and click <b>Threshold Violations</b> .
Alerts page	To view the violations of performance thresholds for all resources, go to <b>Home &gt; Alerts</b> .

In the stand-alone GUI, you can view information about threshold violations in Constraint Violation reports. To view alerts that are generated by threshold violations, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert logs** in the navigation tree and click **All**.

**Tip:** If a performance monitor is already collecting data about a resource when you add, modify, or remove a threshold for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the threshold is updated.

### Defining a performance threshold alert for multiple resources

Performance alerts are triggered by violations of performance thresholds on resources. To specify the performance thresholds, you must edit the performance alert definition.

When you define a performance threshold, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, a threshold violation occurs. For example, you can define a performance threshold that alerts you when the overall back-end response times for managed disks on a SAN Volume Controller exceed a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.



**Tip:** Thresholds are checked against all the internal resources of a specific type; you cannot apply different thresholds to internal resources of the same type. For example, the storage system has two ports, Port A and Port B. You cannot set

different threshold values for each port on the storage system. You must set the same threshold values for all of the ports on the storage system.

1. To select resources, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click Storage Systems.	Choose this option to select a storage system.
From the Network menu, click Switches.	Choose this option to select a switch.

2. To select multiple resources, press Ctrl and click each resource.
3. Click **Actions > Edit Alert Definitions**.
4. From the category list, click **Performance**.
5. To define a performance alert, choose one of the following options:

Option	Description
Click the switch icon  .	Use this option to enable the alert and set the boundary values.
Click the performance chart icon  .	Use this option to edit the boundary values for an enabled alert. <b>Tip:</b> Hover the mouse over the performance chart icon to see the boundary values for an enabled alert.

6. Enter boundary values for triggering warning performance alerts. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than 500 ops/s or less than 200 ops/s, in the **Warning > Greater than** field, enter the value 500. In the **Warning > Less than** field, enter the value 200.
7. Enter boundary values for triggering critical performance alerts. For example, to trigger an alert if the Total I/O Rate for a storage system is greater than 600 ops/s and less than 100 ops/s, in the **Critical > Greater than** field, enter the value 600. In the **Critical > Less than** field, enter the value 100.

#### Tips:

- Tivoli Storage Productivity Center provides recommended values for thresholds that do not vary much between environments. For example, the default critical and stress boundaries for the disk utilization threshold are 80% and 50%. However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine boundary values for a resource, collect performance data over a period of time to establish a base-line of the normal and expected performance behavior for that resource. After you determine a set of base-line values, define thresholds to trigger if the measured performance behavior falls outside the normally expected range.
  - For some metrics and performance thresholds, lower values indicate more stress and higher values indicate idle behavior. For example, a lower value in the critical stress boundary for the Cache Holding Time Threshold might indicate a performance problem.
  - If you do not want to monitor violations for certain boundaries, leave those boundaries blank.
8. To evaluate performance information for the resource, click **Top 10** or **Bottom 10** to view the chart for resources ordered by performance, click a time period,

or use the date and time fields to customize the performance chart display. In the chart, the yellow horizontal lines represent the values that you enter in the **Warning** fields. The red horizontal lines represent the values that you enter in the **Critical** fields.

**Tip:** You can use the information from the performance chart to help you decide the boundary values to set for the alert.

9. To filter any boundary violations when a performance threshold for a related alert is reached, select the check box to ignore that condition. Specify a threshold value for the related alert.
10. Click **Done**.
11. Click **Save**.

Tivoli Storage Productivity Center checks the value for a metric at each collection interval of a performance monitor to determine if a threshold is violated. If a violation occurs, you can view information about that violation in the following locations in the web-based GUI:

*Table 53. Locations where you can view violation of performance thresholds.*

Location in GUI	Description
Storage Systems page Switches page	To view all the violations of performance thresholds on storage systems, go to <b>Storage &gt; Storage Systems</b> and click <b>Threshold Violations</b> .  To view all the violations of performance thresholds on switches, go to <b>Storage &gt; Switches</b> and click <b>Threshold Violations</b> .
Details page for a specific storage systems Details page for a specific switch	To view all the violations of performance thresholds on a specific storage systems, go to <b>Storage &gt; Storage Systems</b> , double click the storage system that you want to view, and click <b>Alerts</b> on the details page.  To view all the violations of performance thresholds on a specific switches, go to <b>Storage &gt; Switches</b> , double click the switch that you want to view, and click <b>Alerts</b> on the details page.
Performance Monitors page	To view the violations of performance thresholds for all resources, go to <b>Home &gt; Performance Monitors</b> and click <b>Threshold Violations</b> .
Alerts page	To view the violations of performance thresholds for all resources, go to <b>Home &gt; Alerts</b> .

In the stand-alone GUI, you can view information about threshold violations in Constraint Violation reports. To view alerts that are generated by threshold violations, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert logs** in the navigation tree and click **All**.

**Tip:** If a performance monitor is already collecting data about a resource when you add, modify, or remove a threshold for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the threshold is updated.

## Suppressing a performance alert for a resource

To suppress the notifications for performance alerts, you edit the performance alert definition.

1. To select a resource, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click Storage Systems.	Choose this option to suppress notifications for an alert for a storage system.
From the Servers menu, click Servers.	Choose this option to suppress notifications for an alert for a server.
From the Servers menu, click Hypervisors.	Choose this option to suppress notifications for an alert for a hypervisor.
From the Network menu, click Switches.	Choose this option to suppress notifications for an alert for a switch.
From the Network menu, click Fabrics.	Choose this option to suppress notifications for an alert for a fabric.

2. To edit the definition for a specific alert for a resource, choose one of the following options:

Option	Description
Right-click the resource in the list and click Edit Alert Definitions.	Use this option to edit the alert for a resource from the resource list page.
Right-click the resource and click View Details. The resource page is displayed. In the General section, click Alerts, and then click the Definitions tab. Click Edit.	Use this option to edit the alert definition from the resource page.

3. From the category list, click **Performance**.

4. To suppress notifications for a performance alert, click the suppression icon  for the alert.

**Tip:** Hover the mouse over the suppression icon to see the customized suppression settings for an enabled performance alert.


5. Select **Override notification settings**.
6. To suppress an alert if only a single threshold violation occurs in a specified time period, select **Suppress alerts unless thresholds are violated more than once within the following length of time**. Specify the number of minutes.
7. To suppress the second and subsequent alerts when a threshold violation occurs repeatedly during a specified time period, select **Suppress alerts if thresholds are violated within the following length of time after the first occurrence**. Specify the number of minutes.
8. Click **Done**.
9. Click **Save**.

## Suppressing a performance alert for multiple resources

To suppress the notifications for performance alerts for multiple resources, you edit the performance alert definition.

1. To select a resource, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click Storage Systems.	Choose this option to suppress notifications for an alert for storage systems.
From the Servers menu, click Servers.	Choose this option to suppress notifications for an alert for servers.
From the Servers menu, click Hypervisors.	Choose this option to suppress notifications for an alert for hypervisors.
From the Network menu, click Switches.	Choose this option to suppress notifications for an alert for switches.
From the Network menu, click Fabrics.	Choose this option to suppress notifications for an alert for fabrics.

- To select multiple resources, press Ctrl and right-click each resource. Click **Edit Alert Definitions**.
- From the category list, click **Performance**.
- To suppress notifications for a performance alert, click the suppression icon  for the alert.

**Tip:** Hover the mouse over the suppression icon to see the customized suppression settings for an enabled performance alert.

- Select **Override notification settings**.
- To suppress an alert if only a single threshold violation occurs in a specified time period, select **Suppress alerts unless thresholds are violated more than once within the following length of time**. Specify the number of minutes.
- To suppress the second and subsequent alerts when a threshold violation occurs repeatedly during a specified time period, select **Suppress alerts if thresholds are violated within the following length of time after the first occurrence**. Specify the number of minutes.
- Click **Done**.
- Click **Save**.

## Viewing alert definitions

View the conditions that trigger alerts for resources and the notification settings for those conditions. You can view the alert definitions for a single resource or for multiple resources.

### Viewing alert definitions for a resource

You can view the alert definitions for a resource to see the triggering conditions, the settings for scripts and notifications, and the threshold values for certain alerts.

- To select a resource, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click Storage Systems.	Choose this option to view alerts for a storage system.
From the Servers menu, click Servers.	Choose this option to view alerts for a server.
From the Servers menu, click Hypervisors.	Choose this option to view alerts for a hypervisor.
From the Network menu, click Switches.	Choose this option to view alerts for a switch.




Option	Description
From the Network menu, click <b>Fabrics</b> .	Choose this option to view alerts for a fabric.

- To view the alert definitions for a resource, choose one of the following options:

Option	Description
Right-click the resource in the list and click <b>Edit Alert Definitions</b> .	Use this option to view the alert definitions for a resource from the resource list page.
Right-click the resource and click <b>View Details</b> . The resource page is displayed. In the General section, click <b>Alerts</b> , and then click the <b>Definitions</b> tab. Click <b>Edit</b> .	Use this option to view the alert definitions from the resource page.

- Optional: To view performance alerts or alerts for a resource category, choose one of the following options:


Option	Description
Click <b>Performance</b> from the category list.	Use this option to view the performance alerts for the resource.
Click a resource category from the list.	Use this option to view the alerts for any of the categories that are related to the selected resource. For example, to view an alert for storage system volumes, click <b>Volumes</b> .

- To view the customized script and notification settings for an enabled alert, click the envelope icon . The alert **Settings** window opens.


**Tip:** To see the customized settings for an enabled alert, hover the mouse over the envelope icon.

- View the detailed settings for scripts and email notifications.


Option	Description
<b>Run Script</b>	If a script is set to run when the alert condition is detected, click <b>Change File</b> to view the script information.
<b>Override Notification Settings</b>	View any overrides for the notification settings for the alert. <b>Tip:</b> If email is selected, click <b>Customize</b> to view the email text.

- To close the alert settings window, click **Cancel**.
- To view the values for an alert condition with threshold values, click the boundary values icon .

**Tip:** To see the threshold values for an enabled alert, hover the mouse over the boundary values icon.

- To close the alert condition window, click **Cancel**.
- To view the detailed settings for a performance alert, click the performance chart icon .

**Tip:** To see the boundary values for an enabled alert, hover the mouse over the performance chart icon.

10. To view performance information for the resource, click **Top 10** or **Bottom 10** to view the chart for resources ordered by performance, click a time period, or use the date and time fields to customize the performance chart display.
11. Click **Done**.
12. To view the customized suppression settings for performance alerts, click the suppression icon .




**Tip:** To see the customized suppression settings for an enabled performance alert, hover the mouse over the suppression icon.

13. Click **Done**.
14. On the **Definitions** tab, click **Cancel** to close the **Edit Alert Definitions** window.




### Viewing an alert definition for multiple resources

You can view an alert definition for multiple resources to see the triggering condition, the settings for scripts and notifications, and the threshold values for certain alerts.

When you select multiple resources, the switch icon for each alert condition indicates how the condition is applied to the resources:

Icon	Status of the alert definition
	The alert definition is disabled.
	The alert definition is enabled.
	The alert definition is enabled for some of the selected resources and disabled for the other selected resources.

When you view an alert definition for multiple resources, the check box choice for each notification setting indicates how the setting is applied to the resources:

Icon	Status of the alert notification
	The notification setting applies to all of the selected resources.
	The notification setting does not apply to any of the selected resources.
	The notification setting has different values for the selected resources.

1. To view an alert, choose one of the following options from the menu bar:

Option	Description
<b>From the Storage menu, click Storage Systems.</b>	Choose this option to view alerts for storage systems.
<b>From the Servers menu, click Servers.</b>	Choose this option to view alerts for servers.
<b>From the Servers menu, click Hypervisors.</b>	Choose this option to view alerts for hypervisor.
<b>From the Network menu, click Switches.</b>	Choose this option to view alerts for switches.


Option	Description
From the Network menu, click <b>Fabrics</b> .	Choose this option to view alerts for fabrics.

- To select multiple resources, press Ctrl and right-click each resource. Click **Edit Alert Definitions**. The **Edit Alert Definitions** window opens.

**Tip:** If multiple resources share setting values, you can hover the mouse over an envelope icon to see any overridden settings for an enabled alert.

- Optional: To view performance alerts or alerts for a resource category, choose one of the following options:

Option	Description
Click <b>Performance from the category list</b> .	Use this option to view the performance alerts for the resource.
Click a resource category from the list.	Use this option to view the alerts for any of the categories that are related to the selected resource. For example, to view an alert for storage system volumes, click <b>Volumes</b> .

- To view the customized script and notification settings for an enabled alert, click the envelope icon . The alert **Settings** window opens.

**Tip:** To see the customized settings for an enabled alert, hover the mouse over the envelope icon.


- View the detailed settings for scripts and email notifications.

Option	Description
<b>Run Script</b>	If a script is set to run when the alert condition is detected, click <b>Change File</b> to view the script information.
<b>Override Notification Settings</b>	View any overrides for the notification settings for the alert. <b>Tip:</b> If email is selected, click <b>Customize</b> to view the email text.

- To close the alert settings window, click **Cancel**.
- To view the values for an alert condition with threshold values, click the

boundary values icon .

**Tip:** If any of the threshold values are different for the selected resources, the threshold fields are blank.

- Click **Cancel** to close the alert condition window.
- To view the detailed settings for a performance alert, click the performance chart icon .

**Tip:** To see the boundary values for an enabled alert, hover the mouse over the performance chart icon.

- To view performance information for the resource, click **Top 10** or **Bottom 10** to view the chart for resources ordered by performance, click a time period, or use the date and time fields to customize the performance chart display.
- Click **Done**.

12. To view suppression settings for performance alerts, click the suppression icon



**Tip:** To see the customized suppression settings for an enabled performance alert, hover the mouse over the suppression icon.

13. Click **Done**.
14. On the **Definitions** tab, click **Cancel** to close the **Edit Alert Definitions** window.

## Viewing and modifying alert notification settings

You can view and modify the alert notification settings to determine the actions that are taken when alert conditions are detected for resources.

The settings are applied to all of the conditions that are specified for the resources. You can view and modify the alert notification settings for a single resource or for multiple resources.

### Viewing the alert notification settings for a resource

You can view the notification settings that determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the conditions that are enabled for the resource.

1. To select the type of top-level resource, choose one of the following options from the menu bar:

Option	Description
From the <b>Storage</b> menu, click <b>Storage Systems</b> .	Choose this option to select a storage system.
From the <b>Servers</b> menu, click <b>Servers</b> .	Choose this option to select a server.
From the <b>Servers</b> menu, click <b>Hypervisors</b> .	Choose this option to select a hypervisor.
From the <b>Network</b> menu, click <b>Switches</b> .	Choose this option to select a switch.
From the <b>Network</b> menu, click <b>Fabrics</b> .	Choose this option to select a fabric.

2. To view the alert notification settings for a resource, choose one of the following options:

Option	Description
Right-click the resource in the list and click <b>Edit Alert Notification Settings</b> .	Use this option to view the notification settings for an alert from the resource list page.
Right-click the resource and click <b>View Details</b> . The resource page is displayed. In the <b>General</b> section, click <b>Alerts</b> , and then click the <b>Notification Settings</b> tab.	Use this option to view the notification settings for an alert from the resource page.

3. View the notification settings that apply to all of the alert definitions that are enabled for the resource.
4. To close the **Notification Settings** page, click **Cancel**.

### Viewing the alert notification settings for multiple resources

You can view the notification settings that determine the actions that are taken when alert conditions are detected for resources. The settings are applied to all of the conditions that are enabled for the resources.

When you select multiple resources, the check box choice for each notification setting indicates how the setting is applied to the resources:

Icon	Status of the alert notification
<input checked="" type="checkbox"/>	The notification setting applies to all of the selected resources.
<input type="checkbox"/>	The notification setting does not apply to any of the selected resources.
<input type="checkbox"/>	The notification setting has different values for the selected resources.

When you select multiple resources, the check box choice for **Suppress Alerts** indicates how the alert suppression rules are applied to the resources:

Icon	Status of the alert notification
<input checked="" type="checkbox"/>	The <b>Suppress Alerts</b> setting applies to all of the selected resources, the suppression options match, but the numbers of minutes might be the same or different. For example, the option to suppress alerts unless thresholds are violated more than once within a specified length of time is enabled for both a Cisco and a Brocade switch. However, the specified length of time for the Cisco switch is set to 10 minutes, while for the Brocade switch the time period is set 20 minutes.
<input type="checkbox"/>	The <b>Suppress Alerts</b> setting applies to none of the selected resources.
<input type="checkbox"/>	The <b>Suppress Alerts</b> setting applies to some of the selected resources. If you click the check box, the first suppression option is selected, but the minutes field is clear. You can then set the suppression options and minutes.

1. To select the type of top-level resource, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click Storage Systems.	Choose this option to select a storage system.
From the Servers menu, click Servers.	Choose this option to select a server.
From the Servers menu, click Hypervisors.	Choose this option to select a hypervisor.
From the Network menu, click Switches.	Choose this option to select a switch.
From the Network menu, click Fabrics.	Choose this option to select a fabric.

2. To view the alert notification settings for multiple resources, press Ctrl and right-click each resource. Click **Edit Alert Notification Settings**.
3. To close the **Notification Settings** page, click **Cancel**.

## Modifying the alert notification settings for a resource

You can modify the alert notification settings to determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the conditions that are specified for the resource.

Configure the notification settings and the scripts for the alerts.

1. To select the type of top-level resource, choose one of the following options from the menu bar:

Option	Description
From the Storage menu, click Storage Systems.	Choose this option to select a storage system.
From the Servers menu, click Servers.	Choose this option to select a server.
From the Servers menu, click Hypervisors.	Choose this option to select a hypervisor.
From the Network menu, click Switches.	Choose this option to select a switch.
From the Network menu, click Fabrics.	Choose this option to select a fabric.

2. To modify the alert notification settings for a resource, choose one of the following options:

Option	Description
Right-click the resource in the list and click Edit Alert Notification Settings.	Use this option to edit the notification settings for an alert from the resource list page.
Right-click the resource and click View Details. The resource page is displayed. In the General section, click Alerts, and then click the Notification Settings tab. Click Edit.	Use this option to edit the notification settings for an alert from the resource page.

3. Determine how you want to be notified when the configuration of a resource changes.
4. Determine how you want to be notified when a performance threshold is violated.
5. Optional: For performance thresholds, suppress alerts if they meet certain criteria.
6. Click **Save**.

## Modifying the alert notification settings for multiple resources

You can modify the alert notification settings to determine the actions that are taken when alert conditions are detected for resources. The settings are applied to all of the conditions that are specified for the selected resources.

Configure the notification settings and the scripts for the alerts.

When you select multiple resources, the check box choice for each notification setting indicates how the setting is applied to the resources:

Icon	Status of the alert notification
<input checked="" type="checkbox"/>	The notification setting applies to all of the selected resources.
<input type="checkbox"/>	The notification setting does not apply to any of the selected resources.

Icon	Status of the alert notification
<input type="radio"/>	The notification setting has different values for the selected resources.

When you select multiple resources, the check box choice for **Suppress Alerts** indicates how the alert suppression rules are applied to the resources:

Icon	Status of the alert notification
<input checked="" type="checkbox"/>	The <b>Suppress Alerts</b> setting applies to all of the selected resources, the suppression options match, but the numbers of minutes might be the same or different. For example, the option to suppress alerts unless thresholds are violated more than once within a specified length of time is enabled for both a Cisco and a Brocade switch. However, the specified length of time for the Cisco switch is set to 10 minutes, while for the Brocade switch the time period is set 20 minutes.
<input type="checkbox"/>	The <b>Suppress Alerts</b> setting applies to none of the selected resources.
<input type="checkbox"/>	The <b>Suppress Alerts</b> setting applies to some of the selected resources. If you click the check box, the first suppression option is selected, but the minutes field is clear. You can then set the suppression options and minutes.

1. To select the type of top-level resource, choose one of the following options from the menu bar:

Option	Description
<b>From the Storage menu, click Storage Systems.</b>	Choose this option to select storage systems.
<b>From the Servers menu, click Servers.</b>	Choose this option to select servers.
<b>From the Servers menu, click Hypervisors.</b>	Choose this option to select hypervisors.
<b>From the Network menu, click Switches.</b>	Choose this option to select switches.
<b>From the Network menu, click Fabrics.</b>	Choose this option to select fabrics.



2. To modify the alert notification settings for multiple resources, press Ctrl and right-click each resource. Click **Edit Alert Notification Settings**.
3. Determine how you want to be notified when the configuration of the resources changes.
4. Determine how you want to be notified when a performance threshold is violated.
5. Optional: For performance thresholds, suppress alerts if they meet certain criteria.
6. Click **Save**.

## Disabling or enabling alert definitions

Disable an alert definition so that Tivoli Storage Productivity Center does not generate notifications for its specified condition. You can re-enable alert definitions at any time.

In Tivoli Storage Productivity Center v5.2.4 and later, many triggering conditions for alerts were migrated from the stand-alone GUI to the web-based GUI. For the conditions that were not migrated, you must use the stand-alone GUI to disable and enable their alert definitions.

Table 54. Disable or enable an alert definition

Action	Procedure
Disabling an alert definition in the web-based GUI.	<ol style="list-style-type: none"> <li>1. In the menu bar, select the resource that contains the alert definition to disable. For example, to disable the alert definition for a storage system, select <b>Storage &gt; Storage Systems</b>.</li> <li>2. Right-click the resource that contains the alert definition and select <b>Edit Alert Definitions</b>.</li> <li>3. Locate the alert definition that you want to disable.</li> <li>4. Click the switch icon next to the definition so that it appears like this:  . <b>Tip:</b> To enable the definition again, click the icon so that it appears like this:  .</li> <li>5. Click <b>Save</b>.</li> </ol>
Disabling an alert definition in the stand-alone GUI.	<ol style="list-style-type: none"> <li>1. In the navigation tree, expand the node that contains the alert definition to disable: <ul style="list-style-type: none"> <li>• <b>Data Manager&gt;Alerting&gt;Computer Alerts</b></li> <li>• <b>Data Manager&gt;Alerting&gt;Directory Alerts</b></li> <li>• <b>Data Manager&gt;Alerting&gt;Hypervisor Alerts</b></li> <li>• <b>Fabric Manager&gt;Alerting&gt;Fabric Alerts</b></li> <li>• <b>Fabric Manager&gt;Alerting&gt;Endpoint Device Alerts</b></li> <li>• <b>Replication Manager&gt;Alerting&gt;Replication Alerts</b></li> </ul> </li> <li>2. Click the name of alert definition that you want to disable.</li> <li>3. On the Edit Alert page, clear the <b>Enabled</b> check box. <b>Tip:</b> To enable the alert definition again, select the <b>Enabled</b> check box.</li> <li>4. Click <b>File &gt; Save</b>.</li> </ol>

## Modifying alert definitions

Modify an alert definition to change its notification settings, threshold values, and suppression settings.



In Tivoli Storage Productivity Center v5.2.4 and later, many triggering conditions for alerts were migrated from the stand-alone GUI to the web-based GUI. For the conditions that were not migrated, you must use the stand-alone GUI to edit their alert definitions.

Table 55. Modifying an alert definition








Action	Procedure
<p>Modifying an alert definition in the web-based GUI.</p>	<ol style="list-style-type: none"> <li>1. In the menu bar, select the resource that contains the alert definition to edit. For example, to edit the alert definition for a storage system, select <b>Storage &gt; Storage Systems</b>.</li> <li>2. Right-click the resource that contains the alert definition and select <b>Edit Alert Definitions</b>.</li> <li>3. Locate the alert definition that you want to edit.</li> <li>4. To enable or disable an alert definition, click the switch icon: <ul style="list-style-type: none"> <li>•  The alert definition is disabled.</li> <li>•  The alert definition is enabled.</li> <li>•  The alert definition is enabled, but not for all the resources that you selected.</li> </ul> </li> <li>5. To edit the notification settings, boundary values, suppression settings for an alert definition, click the following icons: <ul style="list-style-type: none"> <li>•  Edit the notification settings for a definition.</li> <li>•  Edit the threshold boundaries for the storage space on a resource.</li> <li>•  Edit the threshold boundaries for the performance of a resource.</li> <li>•  Edit the suppression settings for the alert definition that override the global setting for the resource. Use this option when you want custom suppression settings for a specific alert definition.</li> </ul> </li> <li>6. Click <b>Save</b>.</li> </ol>

Table 55. Modifying an alert definition (continued)

Action	Procedure
Modifying an alert definition in the stand-alone GUI.	<ol style="list-style-type: none"> <li>In the navigation tree, expand the node that contains the alert definition to edit: <ul style="list-style-type: none"> <li>Data Manager&gt;Alerting&gt;Computer Alerts</li> <li>Data Manager&gt;Alerting&gt;Directory Alerts</li> <li>Data Manager&gt;Alerting&gt;Hypervisor Alerts</li> <li>Fabric Manager&gt;Alerting&gt;Fabric Alerts</li> <li>Fabric Manager&gt;Alerting&gt;Endpoint Device Alerts</li> <li>Replication Manager&gt;Alerting&gt;Replication Alerts</li> </ul> </li> <li>Click the name of the alert definition that you want to edit.</li> <li>On the Edit Alert page, edit the settings for the alert definition.</li> <li>Click <b>File &gt; Save</b>.</li> </ol>

## Renaming alert definitions

In the stand-alone GUI, you can rename alert definitions at any time. This action might be helpful if the original name is ambiguous or lacks clarity.

In Tivoli Storage Productivity Center v5.2.4 and later, many triggering conditions for alerts were migrated from the stand-alone GUI to the web-based GUI. In the web-based GUI, names are not assigned to alert definitions, so the renaming ability is not needed. Use the following steps only for alert definitions that remain in the stand-alone GUI.

To rename an alert definitions, complete the following steps:

- In the navigation tree of the stand-alone GUI, complete one of the following actions:

Alert type	Action
For computer alerts	Expand <b>Data Manager &gt; Alerting &gt; Computer Alerts</b>
For directory alerts	Expand <b>Data Manager &gt; Alerting &gt; Directory Alerts</b>
For hypervisor alerts	Expand <b>Data Manager &gt; Alerting &gt; Hypervisor Alerts</b>
For instance alerts	Expand <b>Data Manager for Databases &gt; Alerting &gt; Instance Alerts</b>
For database-tablespace alerts	Expand <b>Data Manager for Databases &gt; Alerting &gt; Database-Tablespace Alerts</b>
For table alerts	Expand <b>Data Manager for Databases &gt; Alerting &gt; Table Alerts</b>
For fabric alerts	Expand <b>Fabric Manager &gt; Alerting &gt; Fabric Alerts</b>

Alert type	Action
For endpoint device alerts	Expand <b>Fabric Manager &gt; Alerting &gt; Endpoint Device Alerts</b>
For replication alerts	Expand <b>Replication Manager &gt; Alerting &gt; Replication Alerts</b>

The list of existing alerts is displayed.

2. Right-click the alert, and click **Rename**. The Rename Alert window opens.
3. Type the new name, and click **OK**.

## Viewing all alerts

Use the Alerts page in the web-based GUI as a central location to view and manage all the alerts that were detected by Tivoli Storage Productivity Center.

The Alerts page is organized into two sections:

### Alerts summary chart



The summary chart at the top of the page is a visual representation of the alerts that were detected on monitored resources. The chart consists of stacked bars with the number and severity of current alerts on the y-axis and the resource type on the x-axis. The y-axis labeling scales dynamically according to the highest number of alerts in any of the categories.


### Alerts list


The alerts list shows information about all the alerts that were detected on monitored resources. Information in the list is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected.


The version of Tivoli Storage Productivity Center that detected an alert determines where information about the alert is displayed. Alerts that were generated in version 5.1 or later are displayed on the Alerts page and in the **Alerts** tab on the pages for the related resources. Alerts that were generated in versions of the product before 5.1, also called migrated alerts, are only displayed on the Alerts page.

**Tip:** Depending on how data is configured for a NetApp filer, alerts are shown on different pages in the interface:

- Alerts for the block data in a NetApp filer are shown in the **Alerts** tab on the Storage Systems page and in the list of alerts on the Alerts page.
- Alerts for the file data in a NetApp filer are only shown in the list of alerts on the Alerts page.

1. In the menu bar in the web-based GUI, go to **Home > Alerts**. Information about alerts is shown in the list of alerts.
2. Optional: To view the number of alerts for a specific type of resource, hover the mouse pointer over the related icon in the summary chart. For example, if you want to view the number of Critical, Warning, and Informational alerts that were triggered by hypervisors, hover the mouse pointer over the **Hypervisors** icon .

- Optional: To filter the list of alerts for a specific type of resource, click the icon for that resource in the summary chart. For example, if you only want to view a list of alerts for hypervisors, click the **Hypervisors** icon .

**Tip:** To view descriptions of the icons and columns on the Jobs page, click the **Help** icon  in the upper-right corner of the page.

## Viewing alert logs

In the stand-alone GUI, you can also use the Alert Log to view alerts from a central location.

There are two alert logs that are available in the stand-alone GUI. The type of alert determines where information about the alert is shown:


- For alerts generated by Data Manager for Databases, expand **Data Manager for Databases > Alerting > Alert Log**. You can view:
  - All alerts or alerts organized by RDBMS such as by Oracle, SQL Server, Sybase, and UDB. If you do not have license for a specific RDBMS, no nodes for that RDBMS are displayed under the **Alert Log** node.
  - Alerts by type (such as by Instance, Database-Tablespace, Table, User)
  - Alerts directed to you (the user currently logged on to the UI component)
- For alerts generated by all other components, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert Log**. The alert types that are displayed depend on the Tivoli Storage Productivity Center license type. The following is a list of the possible alert types that are displayed under **Alert Log**:
  - All
  - Alerts Directed to <user-id> - Alerts directed to the user logged on to the interface.
  - Authentication Configuration
  - Replication
  - Storage Subsystem
  - Computer
  - Disk
  - Filesystem
  - Directory
  - User
  - OS User Group
  - Fabric
  - Switch
  - Endpoint Device
  - External
  - Configuration Analysis
  - Hypervisor

For alert types that are not displayed for certain license types, see Chapter 2, "Product licenses," on page 17

Each alert node has an icon next to it that indicates its status:

- Red-outlined circle: a negative alert occurred, such as job failed, new disk defect.
- Green-outlined box: a neutral alert was detected, such as a new disk detected.

When you click an alert node, the alert log window opens a table that contains

information for all the alerts for that node. Click  on a row to drill down for more information.

The following table describes the buttons on the Alert Log window:

Button	Description
Clear	Clear the selected alerts. When you click this button, the icon for the alert changes to indicate that the alert was cleared. The alert continues to be displayed in the alert log.
Delete	Delete the selected alerts from the alert log.
Clear ALL	Clear all of the alerts in the report. When you click this button, the icon for all the alerts changes to indicate that the alerts were cleared. The alerts continue to be displayed in the alert log.
Delete ALL	Delete all of the alerts in the current report from the alert log.
Refresh	Regenerate the report you are viewing to pick up any newly generated alerts that were triggered since you last generated the report.

**Tip:** Some alerts are shown automatically when a client application is registered to receive alerts from Tivoli Storage Productivity Center. You can identify these alerts by the keyword *Pre-defined* in their names. You can hide these alerts from view in the GUI. For more information, see the IBM Tivoli Storage Productivity Center Problem Determination Guide.

**Restriction:** For database tables, the values in the **Resource Type** and **Resource Name** columns are set to N/A for the following alerts:

- Total Table Size Greater Than
- Chained Row Count Greater Than

## Viewing alerts for a type of resource

Use resource list pages in the web-based GUI to view only the alerts for a specific type of top-level resource. Each resource list page has an **Alerts** tab that lists only the alerts for monitored storage systems, servers, hypervisors, switches, or fabrics. For example, you can use the list page for storage systems to view only the alerts that were detected on monitored storage systems.


1. In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want to view the alerts that are related to storage systems, go to **Storage > Storage Systems**.
2. Click the **Alerts** tab to view a list of all the alerts for the monitored resources of that type. The information about alerts is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected.

**Tip:** The icon on the **Alerts** tab represents the most severe alert condition that was detected on the monitored resources. For example, if one of the jobs for storage systems failed, the following icon is shown on the tab:



**Alerts.**

3. Optional: View the icons at the top of the **Alerts** tab for a summary of alert severities. This summary includes the number of alerts that have a Critical, Warning, Informational, or Acknowledged status.

**Tip:** To view descriptions of the icons and columns on the **Alerts** tab, click the **Help** icon  in the upper-right corner of the page.

## Acknowledging alerts

Some alerts in Tivoli Storage Productivity Center are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

**Tip:** Acknowledging alerts does not affect the status of resources. The status of a resource is determined by its consolidated status and the propagated status of its internal resources, not by alerts that are generated. For information about how to acknowledge the status of a resource, see “Acknowledging the status of resources” on page 202.

To acknowledge an alert, complete the following steps:

1. In the web-based GUI, go to a list of alerts that includes the alerts that you want to acknowledge. For example, if you want to acknowledge the alerts for a specific server, complete the following steps.
  - a. In the menu bar in the web-based GUI, go to **Servers > Servers..**
  - b. Right-click the server that has alerts you want to acknowledge and select **View Details**.
  - c. Click the **Alerts** link.
2. Right-click the alert that you want to acknowledge and select **Acknowledge alerts**. The severity of the alert is changed to Acknowledged on the Alerts page, resource list page, and resource details page.

For example, if you acknowledge a server alert that has a Warning severity, it is shown on the Alerts page, Servers page, and Server details page with the following severity:



### **Warning - Acknowledged**

3. Optional: To acknowledge multiple alerts at the same time, press Ctrl and click each alert. To select a series of alerts, select the first alert, and then press Shift and click the last alert. When the alerts are selected, right-click a row in the list and select **Acknowledge alerts**.
4. Optional: To acknowledge all the alerts that are shown in a list, select **Acknowledge all alerts** from the **Actions** menu.
5. Optional: To acknowledge all the alerts with an Informational severity in a list, select **Acknowledged informational alerts** from the **Actions** menu. You might want to acknowledge all Informational alerts when you want to view only alerts that represent potential or serious problems that require action to resolve.

## Removing alerts

Remove alerts when you no longer require to view them.

By default, alerts are automatically removed based on retention settings that are defined on the History Retention page in the web-based GUI. However, you can use the following procedure to manually remove alerts at any time.

1. In the web-based GUI, go to a list of alerts that includes the alerts that you want to remove. For example, if you want to remove alerts for all different

types of top-level resources, go to **Home > Alerts**. If you want to remove alerts for a specific type of resource such as storage systems, go to **Storage > Storage Systems**, and click the **Alerts** tab.

2. Right-click the alert that you want to remove and select **Remove alerts**.
3. Click **Remove** to remove the alert. When you remove an alert, it is no longer shown on the Alerts page or in the **Alerts** tabs on the related resource list and resource detail pages.

For example, if you remove an alert for a fabric, that alert is removed from the Alerts page, Fabrics page, and Fabric details page.

4. Optional: To remove all the alerts that are shown in a list, select **Remove all alerts** from the **Actions** menu.
5. Optional: To remove all the acknowledged alerts that are shown in a list, select **Remove acknowledged alerts** from the **Actions** menu.

---

## Viewing information about resources

You can view detailed information about the resources that are monitored by Tivoli Storage Productivity Center. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

**Remember:** Data collection jobs retrieve information about monitored resources. To ensure that the latest information about resources is available in the web-based GUI, run data collection jobs regularly. To learn about how to collect data, see “Collecting data” on page 145.

## How information is organized

Information about monitored resources is organized into different pages in the web-based GUI. You can use the pages to view summary information about all resources, detailed information about top-level resources, and detailed information about internal and related resources.

### The dashboard

The dashboard is displayed when you start a new session of Tivoli Storage Productivity Center or when you go to **Home > Dashboard** in the menu bar. Use the dashboard to view the following summary information about a storage environment:

- A topology or resource diagram of the top-level resources that are monitored by Tivoli Storage Productivity Center. This diagram shows the number of top-level resources that are monitored and their overall status. You can also click the icons in the resource diagram to view more detailed information about a top-level resource on a resource list page.
- Capacity information that was aggregated for storage systems, servers, and hypervisors. Use this information to view the overall storage capacity for each type of resource. This information is presented as horizontal bar charts.
- The status of Tivoli Storage Productivity Center jobs that were run during a specific time. You can also click links to access a complete list of jobs.
- The number of unacknowledged alerts that were detected on resources during a specific time. You can also click a link to access a complete list of alerts that were detected on all monitored resources.

For more information about the dashboard, see “Dashboard view” on page 188.

## Performance Monitors page

Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Use the Performance Monitors page to view and manage performance monitors and to track threshold violations. Specifically, you can complete the following actions:

- View detailed information about performance monitors. Use this information to identify performance monitors that are encountering problems when they run. You can also view the metrics that are collected by a performance monitor for a resource.
- View detailed information about performance thresholds that were violated. Use this information to identify potential performance problems on monitored resources.

To access the Performance Monitors page, go to **Home > Performance Monitors** in the menu bar.

## Alerts page

Use the Alerts page to view and manage the alert conditions that were detected on monitored resources. Specifically, you can complete the following actions:

- View the overall number of alerts that were detected on monitored resources. This information is presented in a bar chart for the following types of alerts:
  - Database
  - Fabric
  - Hypervisor
  - Other
  - Performance
  - Replication
  - Server
  - Storage system
  - Tivoli Storage Productivity Center
- View information about specific alerts, including severity, last occurrence, and the resources where they were detected.
- Acknowledge alerts that were reviewed but are not yet resolved.
- Unacknowledge alerts.
- Remove alerts from the page.

To access the Alerts page, go to **Home > Alerts** in the menu bar.

## Tasks page

Use the Tasks page to view and manage the tasks that Tivoli Storage Productivity Center uses to optimize resources and provision storage. Tasks are created when you complete any of the following actions:

- Assign storage to servers, hypervisors, or a cluster by using the Provision Storage wizard.
- Move volumes to higher or lower tiers by using the Analyze Tiering wizard.
- Move or convert volumes by using the Transform Storage wizard.
- Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard.



To access the Tasks page, go to **Home > Tasks** in the menu bar.

## System Management page

Use the System Management page to view information about the overall condition of Tivoli Storage Productivity Center. You can view information about the servers on which Tivoli Storage Productivity Center is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.

To access the System Management page, go to **Home > System Management** in the menu bar.

## Resource list pages

Use a resource list page as central location for viewing information about the top-level resources of a specific type. For example, use the Storage Systems page to view a list of the storage systems that are monitored by Tivoli Storage Productivity Center. Each resource list page corresponds to a specific type of resource, which includes storage systems, volumes, pools, servers, hypervisors, switches, and fabrics. You can view the following information on resource list pages:

- The overall status of the resources
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

To learn about how to access the list pages for resources, see “Viewing information about top-level resources” on page 193.

**Note:** Volumes and pools are internal resources of storage systems, but have their own resource list pages so that you can view them all in a central location.

## Resource detail pages

Use resource details pages to view information about resources that are internal or related to a top-level resource. For example, use the Storage System details page to view detailed information about a specific storage system. You can view the following information on resource detail list pages:

- The overall capacity and storage usage for a resource
- Key information about a resource
- Alerts that were generated for a resource
- Jobs that were run for a resource
- Information about components that exist in top-level resources
- Information about the resources that are related to another resource

To learn about how to access detail pages for resources, see “Viewing information about internal and related resources” on page 194.

## Reporting

Use the IBM Tivoli Common Reporting tool in the web-based GUI to view detailed reports about resources. The following reports are available:

### Predefined reports

These reports are included with Tivoli Storage Productivity Center. The following types of predefined reports are available:

- Resources
- Performance
- Historical space
- Storage tiering

### Custom reports

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To learn about how to use the IBM Tivoli Common Reporting tool, see Chapter 7, “Reporting,” on page 483.

## Dashboard view

The dashboard is shown when you first log in to the GUI or when you go to **Home > Dashboard** from the menu bar. Use the dashboard to monitor the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about the status and capacity of resources, the jobs that collect performance data, and the unacknowledged alert conditions that are detected on resources.

Specifically, the dashboard provides a concise, yet detailed overview of your storage environment that includes the following information:

- The status of monitored resources on the resource diagram
- The overall capacity of monitored resources
- The status of performance monitors
- The status of alerts that were detected on monitored resources

**Important:** The information that is shown on the dashboard is gathered by data collection jobs. To ensure that the latest information is available, run these jobs regularly.

### Viewing the status of monitored resources on the resource diagram

The diagram in the top section of the dashboard provides a visual representation of the resource types that are monitored by Tivoli Storage Productivity Center. Use this diagram to view the following information:

#### A list of resources that are being monitored for each resource type

Each type of resource in the diagram is represented by an icon. To learn more detailed about the individual resources that are monitored for a type, click the related icon in the diagram. For example, if you click the icon for storage systems, a new page shows detailed information about the individual storage systems that are being monitored.

#### The number of resources that are monitored for a resource type

A number in parentheses is shown next to each resource icon. This number represents the number of resources for each type that is monitored. For

example, if 20 storage systems are monitored, the value **Storage Systems (20)** is shown. This number excludes resources that are not monitored by Tivoli Storage Productivity Center.

**Tip:** If no resources are being monitored for a resource type, click its icon on the dashboard to add a resource of that type. For example, if the dashboard shows **Storage Systems (0)**, you can click the icon for storage systems to add a storage system. The GUI guides you through the process for adding the resource for monitoring. For more information about adding resources, see “Adding resources” on page 127.

### The overall status for each resource type

A status symbol is shown next to each resource icon. This symbol represents the most critical status that was detected on the internal resources for a resource type. For example, if 20 storage systems are monitored, and an error was detected on a volume for one of those storage systems, a red symbol is shown next to the storage systems icon in the diagram. If no errors, warnings, or unreachable statuses were detected on the internal resources of monitored storage systems, then a green symbol is shown.

Use the status to quickly determine the condition of your storage systems and if any actions must be taken. On the dashboard, the following statuses might be shown for a resource type:



#### Error (red)

An error status was detected on one or more of the monitored resources for a resource type. Error statuses represent serious problems on a resource. Resolve these problems as soon as possible. View the details of a resource to learn more about the status of its internal resources.



#### Unreachable (orange)

One or more of the monitored resources for a resource type are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.



#### Warning (yellow)

A warning status was detected on one or more of the monitored resources for a resource type. Warning conditions are not critical, but represent potential problems. View the details of a resource to learn more about the status of its internal resources.



#### Normal (green)

No warnings or errors were detected on the monitored resources for a resource type.



#### Unknown status (gray)

One or more of the monitored resources for a resource type have an unknown status. This status might occur if no data was collected about a resource. To change an unknown status, use a data collection job to collect status information about the resource.

For information about how status is determined, see “How status is determined” on page 197.

**The number of statuses that were detected and acknowledged for each resource type** To view the total number of statuses for a resource type, move the mouse pointer over the related status symbol. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a Tivoli Storage Productivity Center user. An acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.

## Viewing the overall capacity of monitored resources

The Capacity section of the dashboard contains a visual summary of the storage capacity for each type of resource that Tivoli Storage Productivity Center monitors.

To cycle through the capacity information for different types of resources, click the arrows ◀ ▶ next to the name of the resource type. Use the chart information to quickly determine the overall storage capacity for monitored resources and their available space. Hover over areas in the charts that are not labeled for more information.

## Storage Systems

Storage systems can be configured with block storage, file storage, or a combination of block and file storage. The type of storage systems that are monitored by Tivoli Storage Productivity Center determines the capacity charts that are shown on the dashboard:

- If only storage systems with block data are monitored, then a single chart is shown with the following information: pool space, volume capacity, available pool space, and volume space that is assigned and unassigned.
- If only storage systems with file data are monitored, then a single chart is shown with the following information: file system capacity, used file system space, and available file system space.
- If storage systems with block data and file data are monitored, then charts for both block and file data are shown.

Use the information in the following horizontal bar charts to view details about the capacity of all the monitored storage systems in your environment:

### Storage systems that are configured with block data

The following capacity information is shown for monitored storage systems that are configured with block data:

#### Pool Space

Total amount of storage space in the pools that are associated with monitored storage systems. For SAN Volume Controller, this value represents the amount of managed disk (MDisk) space that is assigned to managed disk groups. For XIV systems, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

**Tip:** When the managed disks (MDisks) of a storage virtualizer are thin-provisioned on a back-end storage system, the value for **Pool Space** includes the virtual space of those managed disks. This virtual space might not be fully allocated (or allocatable) on the back-end storage systems.

## **Volumes**

Total amount of unique storage system volume space or virtual storage volume space on monitored storage systems. For thin-provisioned volumes, this value is the physical capacity of the volume. In an XIV system, this value is referred to as the real space, or physical ("hard") size, and not the potentially much larger virtual ("soft") size, that is allocated. For NetApp devices, this value represents the total amount of storage space on the Fibre Channel (FC) LUNs that are associated with the devices.

### **Assigned (Volumes)**

Total volume space within monitored storage systems that is mapped or assigned to host systems.

### **Unassigned (Volumes)**

Total volume space within monitored storage systems that is not mapped or assigned to host systems.

### **Reserved**

Amount of pool space that is reserved by provisioning tasks. Pool space is reserved when a provisioning task is created, and assigned when a provisioning task is run.

### **Unreserved**

Amount of space in storage system pools that is not allocated for volumes, and is not reserved by pending or scheduled provisioning tasks.

## **Storage systems that are configured with file data**

The following capacity information is shown for monitored storage systems that are configured with file data:

### **File System Capacity**

Total amount of storage space on all the file systems on the storage systems and filers.

### **Used File System Space**

Amount of file system space on the storage systems and filers that is being used by files and directories.

### **Available File System Space**

Amount of unused storage space on all the file systems on the storage systems and filers.

## **Servers**

The following capacity information is shown for the monitored servers in your environment: .

### **Total Disk**

The total disk space for all the local and SAN-attached storage (from storage systems) on monitored servers.

### **Storage Systems**

The amount of storage space that is assigned to servers from storage systems. The value is only available when SAN-attached storage is assigned to servers.

### **Other (storage systems)**

The amount of storage space on the local disks and other non-SAN-attached disks in servers.

**Used Disk, Available Disk**

The total amount of used and unused disk space on the local and SAN-attached storage for servers.

**FS/LV (File Systems and Logical Volumes), Available FS (File System Space)**

The amount of used and unused space on the file systems and logical volumes for servers.

**Hypervisors**

The following capacity information is shown for the monitored hypervisors in your environment:

**Total Disk**

The total disk space for all the local and SAN-attached storage (from storage systems) on hypervisors.

**Storage Systems**

The amount of storage space that is assigned to hypervisors from storage systems. The value is only available when SAN-attached storage is assigned to hypervisors.

**Other (storage systems)**

The amount of storage space on the local disks in the hypervisor.

**Used Disk, Available Disk**

The amount of used and unused disk space on the local and SAN-attached storage for the hypervisor.

**FS (file system), Direct (direct-attached storage)**

The proportion of file system space to direct-attached storage space that is used on the local and SAN-attached storage for the hypervisor.

Direct-attached storage is storage that not managed by the hypervisor, but is assigned directly to a virtual machine on the hypervisor.

**FS/LV (File Systems and Logical Volumes), Available FS (File System Space)**

The amount of storage space on file systems and logical volumes that is used and available.

**Viewing the status of performance monitors**

A performance monitor is a data collection job that gathers performance data about monitored storage systems and switches. The Performance section of the dashboard shows a status summary of the performance monitors that are active or were run in your environment. This summary includes the number of performance monitors that have the following statuses: Running, Running with problems, Failed, and Not running.



To view detailed information about all performance monitors, click **View all monitors**.

**Viewing the status of alerts that were detected on monitored resources**

Alerts are generated when Tivoli Storage Productivity Center detects certain conditions or events on monitored resources. The Unacknowledged Alerts section of the dashboard shows a status summary of the unacknowledged alerts that were detected in your environment during a specified time period. This summary includes the number of alerts that have the following statuses: Critical, Warning,

and Informational. An unacknowledged alert is an alert that was not reviewed or resolved by a Tivoli Storage Productivity Center user.

Use the following actions to manage the display of statuses for unacknowledged alerts:

- To select the time period in which you want to view alert statuses, click the arrows   to cycle through the choices. You can view the status of all unacknowledged alerts that were detected in the last hour, last day, or last week.
- To view more detailed information about alerts, click **View all alerts** to access the Alerts page.


## Viewing information about top-level resources

Use a resource list page as a central location for viewing information about all the monitored resources of a specific type. Resources that have a list page include storage systems, volumes, pools, shares, servers, hypervisors, switches, and fabrics.

1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view. You can access the following types of resources:
  - **Storage:** storage systems, volumes, pools, shares
  - **Servers:** servers, hypervisors
  - **Network:** switches, fabrics
2. Select the type of resource that you want to view. For example, if you want to view information about servers, go to **Servers > Servers**. A list of monitored servers is displayed on Servers page.
3. Locate the resource that you want to view in the list. Each column in the list shows information about the resource.
4. Optional: To view key attributes of a specific top-level resource without leaving the page, right-click the resource and select **View Properties**. A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
5. Optional: To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**. A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, and related resources.

A resource list page shows information about the monitored resources for the selected type. You can view the following information on this page:

- The overall status of the resources for a specific type
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

**Tip:** To view descriptions of the information that is available on a resource list page, click the **Help** icon  in the upper-right corner of the page.

**Related reference:**

“Resources that you can monitor in the web-based GUI” on page 128

You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

## Viewing information about internal and related resources

Use resource details pages to view information about a specific top-level resource and its internal and related resources.


1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view. You can access the following types of resources:
  - **Storage:** storage systems, volumes, pools
  - **Servers:** servers, hypervisors
  - **Network:** switches, fabrics
2. Select the type of the resource that you want to view. A resource list page shows the monitored resources for the selected type. For example, if you want to view a storage system, go to **Storage > Storage Systems**. A list of monitored storage systems is displayed on the Storage Systems page.
3. In the list of resources, right-click the resource that you want to view and select **View Details**. A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, and related resources.
4. To view information about the internal resources of the top-level resource, click the name of the internal resource in the **Internal Resources** section. A list of the internal resources is displayed. Each column in the list shows information about the resources. For example, on the details page for a server, click **Controllers** to view a list of the controllers that are associated with the server.
5. Optional: To view key attributes of a specific internal resource without leaving the page, right-click the resource and select **View Properties**.
6. To view information about resources that are related to the top-level resource, click the name of the related resource in the **Related Resources** section. A list of the related resources is displayed. For example, on the server details page, click **Storage Systems** to view a list of the storage systems that have at least one disk (volume) assigned to the server.
7. Optional: To view key attributes of a specific related resource without leaving the page, right-click the resource and select **View Properties**.
8. Optional: To view the resource details page for a related resource, right-click the resource and select **View Details**.

A resource details page shows information about a top-level resource. The information is available in the following links and sections on the page:

- **Overview:** The overall capacity and storage usage for a resource.
- **Properties:** Key information about a resource.
- **Alerts:** Alerts that were generated for a resource.
- **Tasks:** Jobs that were run for a resource.
- **Internal Resources:** Information about the components that exist in a higher-level resource.



- **Related Resources:** Information about the resources that are related to another resource. A related resource is external to a higher-level resource, but is related to it through assigned storage, a network connection, or virtual hosting.

**Tip:** To view descriptions of the information that is available on a resource details page, click the **Help** icon  in the upper-right corner of the page.

**Related reference:**

“Resources that you can monitor in the web-based GUI” on page 128  
 You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

## Display differences in the interfaces

Tivoli Storage Productivity Center stores information about monitored resources in a database repository. You can view this information in the stand-alone GUI or in the web-based GUI. Depending on the interface that you use, there might be some differences in how the information is presented.

### The web-based GUI does not show a fabric for a physical, Cisco SAN

The stand-alone GUI shows a "fabric" for a physical, Cisco SAN, whereas the web-based GUI does not show a "fabric" for a physical, Cisco SAN.

- In the stand-alone GUI, Tivoli Storage Productivity Center automatically generates a worldwide name (WWN) to represent a physical, CISCO SAN as a fabric. This WWN typically starts with the numbers "3000". The WWN for the fabric is shown on the **Topology > Fabrics** panel and in reports for fabrics.
- In the web-based GUI, a fabric is not shown for physical, Cisco SANs on the Fabrics page. Only the virtual fabrics, physical switches, and logical switches for Cisco Fibre Channel switches are shown in the web-based GUI.

### The number of unidentified endpoints on fabrics is different from the stand-alone GUI

The web-based GUI and stand-alone GUI might display a different number of unidentified endpoints on a fabric:

- In the web-based GUI, unidentified endpoints are labeled as **Unknown Resources** and each endpoint *port* is considered one **Unknown Resource**.
- In the stand-alone GUI, unidentified endpoints are labeled as **Other** and each endpoint *node* is considered one **Other**.

For example, an endpoint with two ports is shown as one **Other** resource in the stand-alone GUI, but in the web-based GUI it is shown as two **Unknown Resources**.

### Pool space for NetApp filers

For NetApp filers that are configured with block data, the values for pool space that are shown in the web-based GUI are different from the values that are shown in a storage management tool for NetApp filers, such as FilerView. The values that are different include total pool capacity and available pool space.

This discrepancy occurs because the pool space information that is collected by SMI-S agents for Tivoli Storage Productivity Center includes reserved space for the Write Anywhere File Layout (WAFL) file system and the Snapshot session type. Tools such as FilerView do not include this additional space.

### Incorrect values for storage systems in the web-based GUI after upgrade

When you upgrade Tivoli Storage Productivity Center from Version 4.2.x to Version 5.2, some columns on the Storage Systems page might display incorrect values. For example, the Raw Disk Capacity column, the Disks column, and the Volumes column might display a value of 0. To resolve this issue, you must run a probe for the resource when the upgrade completes. When the probe successfully completes, the columns display the correct values.

---

## Monitoring the status of resources

You can use the web-based GUI to monitor the status of resources and identify potential problem areas in a storage environment. Monitored resources include top-level resources and their internal and related resources.

Use status information to quickly determine the condition of your storage and if any actions must be taken. Tivoli Storage Productivity Center provides a number of different statuses to help you determine the condition of resources. Each status is represented by an icon in the web-based GUI.

*Table 56. Possible statuses of monitored resources*









Icon	Status	Description
	Error	A serious problem was detected on a resource or on its internal resources. Resolve these problems as soon as possible.
	Error - Acknowledged	An Error status was detected and acknowledged. An Error - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.  An acknowledged status is not used when determining the status of related, higher-level resources. For example, if the status of a volume is Error, the status of the related storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall status of the storage system.
	Unreachable	A resource is not responding to requests from the Tivoli Storage Productivity Center server. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.
	Unreachable - Acknowledged	An Unreachable status was detected and acknowledged. An Unreachable - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.  An acknowledged status is not used when determining the status of related, higher-level resources. For example, if the status of a controller is Unreachable, the status of the related server is also unreachable. If the Unreachable status of the controller is acknowledged, its status is not used to determine the overall status of the server.

Table 56. Possible statuses of monitored resources (continued)

Icon	Status	Description
	Warning	A Warning status represents potential problems on a resource or on its internal resources. This status is not critical.
	Warning - Acknowledged	<p>A Warning status was detected and acknowledged. An Warning - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>An Warning - Acknowledged status is not used when determining the status of related, higher-level resources. For example, if the status of a controller is Warning, the status of the related server is also Warning. If the Warning status of the controller is acknowledged, its status is not used to determine the overall status of the server.</p>
	Normal	No warnings or errors were detected on a monitored resource.
	Unknown	A resource is known to Tivoli Storage Productivity Center but is not monitored. To change an Unknown status, use a data collection job to collect status information about the resource.

For top-level resources, use the resource details page to view the status of their internal resources.

## How status is determined

Top-level resources that are monitored in the web-based GUI include storage systems, servers, hypervisors, fabrics, and switches. The status of a top-level resource is determined by its consolidated status and the propagated status of its internal resources.

A consolidated status is the status that is reported by the top-level resource for its own hardware. For example, if the cooling fans in a storage system are stopped and the internal temperature is too high, that status is reported to Tivoli Storage Productivity Center.

A propagated status is the combined statuses of the monitored, internal resources for a top-level resource. For example, if a Warning status is detected on a volume in a storage system, that Warning status is used to determine the overall status of the storage system.

Tivoli Storage Productivity Center uses the values for consolidated status and propagated status to calculate the overall status of a top-level resource. The statuses of the following internal resources are used to calculate the status of top-level resources.

Table 57. Internal resources that are used to determine the status of top-level resources




















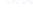











Top-level resource	Internal resources that are used to determine the status of a top-level resource
Fabric 	 ,  Blades, and ports on blades  Ports that are not blades  Switches

Table 57. Internal resources that are used to determine the status of top-level resources (continued)

Top-level resource	Internal resources that are used to determine the status of a top-level resource
<p>Hypervisor</p> 	<ul style="list-style-type: none"> <li> Controllers</li> <li> Ports on controllers</li> <li> Virtual machines</li> </ul>
<p>Server</p> 	<ul style="list-style-type: none"> <li> ,  Controllers, and ports on controllers</li> <li> Volume groups</li> <li> Monitored directories</li> </ul>
<p>Storage system</p> 	<ul style="list-style-type: none"> <li> Clusters</li> <li> Controllers</li> <li> Disks</li> <li> External disks</li> <li> Modules</li> <li> Network shared disks</li> <li> Nodes</li> <li> Pools, and the internal resources that are associated with pools: <ul style="list-style-type: none"> <li> Volumes</li> <li> RAID arrays</li> <li> Managed disks</li> </ul> </li> <li> Ports</li> </ul>
<p>Switch</p> 	<ul style="list-style-type: none"> <li> Blades</li> <li> Logical switches</li> <li> Ports</li> </ul>

Tivoli Storage Productivity Center uses the following statuses of internal resources to help calculate the status top-level resources:

- Unknown
- Normal

- Error
- Warning
- Unreachable

**Tip:** Statuses that are acknowledged are not used to calculate the overall status of a top-level resource.

Internal resources for a top-level resource might have different statuses. Tivoli Storage Productivity Center uses the most severe status of an internal resource when determining the overall status for a top-level resource. For example, in a storage system, a port might have an Error status, a pool might have a Warning status, and a number of controllers might have an Unknown status. In this case, the overall status for the storage system is Error because it is the most severe status that was detected on the internal resources.

Statuses for internal resources are propagated in the following way to determine the status of a top-level resource:

- If an internal resource has an Error status, then the related top-level resource has an Error status.
- If an internal resource has an Unreachable status, and no other internal resources have an Error status, the related top-level resource has an Unreachable status.
- If an internal resource has a Warning status, and no other internal resources have an Error or Unreachable status, the related top-level resource has a Warning status.
- If an internal resource has a Normal status, and no other internal resources have an Error, Unreachable, or Warning status, the related top-level resources has a Normal status.
- If an internal resource has an Unknown status, and no other internal resources have an Error, Unreachable, Warning, or Normal status, the related top-level resource has an Unknown status.

The following table shows some of the possible combinations of statuses and their resulting, propagated status.

*Table 58. Propagation of the statuses for internal resources in the web-based GUI*


























Error 	Unreachable 	Warning 	Normal 	Unknown 	Resulting status for top-level resource
				X	 Unknown
			X		 Normal
			X	X	 Normal
		X			 Warning
		X		X	 Warning
		X	X	X	 Warning
	X				 Unreachable
	X			X	 Unreachable

Table 58. Propagation of the statuses for internal resources in the web-based GUI (continued)

Error 	Unreachable 	Warning 	Normal 	Unknown 	Resulting status for top-level resource
	X		X	X	 Unreachable
	X	X	X	X	 Unreachable
X					 Error
X				X	 Error
X			X	X	 Error
X		X	X	X	 Error
X	X	X	X	X	 Error

## Viewing the overall status of resources

Use the dashboard to monitor the overall status of monitored resources and identify potential problem areas in a storage environment. The dashboard shows the overall status of storage systems, servers, hypervisors, fabrics, and switches, and the total number of different statuses that were detected on the individual resources for each type.

1. In the menu bar, go to **Home > Dashboard**. Information about the overall status of top-level resources is shown on the resource diagram.
2. View the status icon that is shown with each resource icon to determine the overall status of a resource type. The status icon represents the most critical status that was detected on the resources for a specific type.

For example, if Tivoli Storage Productivity Center monitors 20 storage systems, and an error condition was detected for one of those storage systems, the following status icon is shown with the storage systems icon:

 **Storage Systems (20)**

If no Error, Warning, or Unreachable statuses were detected on the monitored storage systems, the following status icon is shown with the storage systems icon:

 **Storage Systems (20)**

3. Move the mouse pointer over the status icon for a resource type to view the number of statuses that were detected and acknowledged for that resource type. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a Tivoli Storage Productivity Center user.

For example, move the mouse pointer over the status symbol for fabrics to view the number of monitored fabrics that have an Error, Unreachable, Warning, Normal, or Unknown status. If five fabrics have an Unreachable status, three fabrics have a Normal status, and no fabrics have an Error or Warning status, the following information is shown when you hover over the status icon for fabrics:



0 Error



5 Unreachable



0 Warning



3 Normal

**Related reference:**

“How status is determined” on page 197

Top-level resources that are monitored in the web-based GUI include storage systems, servers, hypervisors, fabrics, and switches. The status of a top-level resource is determined by its consolidated status and the propagated status of its internal resources.

## Viewing the status of top-level resources

Use a resource list page as a central location for viewing the propagated statuses of top-level resources. Resources that have a list page with status information include storage systems, volumes, pools, servers, hypervisors, switches, and fabrics.

1. In the menu bar, go to the type of top-level resource that you want to view. For example, if you want to view the status of monitored switches, go to **Network > Switches**.

2. To view the total number of different statuses that were detected for the type of resource, check the status icons at the top of the resource list page. The number next to the icon represents how many of each status was detected.

For example, if 20 switches are monitored, and 10 have a Normal status while the other 10 have a Warning status, the following status information is displayed at the top of the page:



10 Normal




10 Warning




0 Error

3. To view the status for specific resources, check the Status column in the list of resources. The value in the Status column represents the propagated status for a top-level resource.

4. Optional: Click the **Search** icon  to filter the list so that only resources with a specified status are shown. You can use filtering to identify the resources with the most serious Error conditions in your environment.

5. Optional: Select **Status**.

6. Optional: In the **Choose Value** list, select the status on which you want to filter. The list of resources is automatically refreshed and only resources with the

filtered status are shown. For example, select **Warning**  to view only the resources on which a Warning status was detected.

## Viewing the status of internal resources

Use a resource details page to view the propagated status of a top-level resource and the status of its internal resources.

1. In the menu bar, go to a type of top-level resource. For example, if you want to view the status of monitored storage systems and their internal resources, go to **Storage > Storage Systems**.

2. Locate the resource in the list and check the Status column. The value in the Status column represents the propagated status for the top-level resource.
3. To view the status of internal resources, right-click the top-level resource and select **View Details**. A resource details page shows the internal resources for the selected resource.
4. Check for any status icons that are associated with an internal resource. If Tivoli Storage Productivity Center detects a status other than normal, an icon is displayed for the most critical status.

For example, if a server has three controllers with an Error status and two controllers with a Warning status, the Error status icon is shown for controllers in the **Internal Resources** section:

 **Controllers (5)** .

5. To view the status of specific internal resources, click the name of the internal resource in the **Internal Resources** section. A list of the monitored internal resources is displayed. For example, on the details page for a server, click **Controllers (5)** to view a list of the five controllers that are associated with the server.
6. Check the Status column to view the status for each internal resource.

**Tip:** When viewing information about some internal resources, the Status column might show values that are more specific than Error, Unreachable, Warning, Unreachable, Unknown, and Normal. Use this additional status to determine the cause of the problem for an internal resource.

For example, in the **Internal Resources** section on the details page for a storage system, an Error status is shown for managed disks:

 **Managed Disks (5)** .

This status means that one or more managed disks have an Error condition. When you view the list of managed disks, the Status column might show an Offline status for a specific managed disk. When an Offline status is detected, the web-based GUI shows it as an Error status on the details page for the storage system.

## Acknowledging the status of resources

The web-based GUI shows Error, Unreachable, and Warning statuses for top-level resources and internal resources. Sometimes these statuses might represent problems that commonly occur but can be ignored. In such cases, you can acknowledge the resources so that they are not used when propagating statuses for top-level resources or when calculating the total number of Error, Warning, and Unreachable conditions that were detected.

### Acknowledging the status of top-level resources

The web-based GUI shows a summary of the statuses for top-level resources on the dashboard and resource list pages. You can acknowledge the status of a top-level resource so that its status is not shown as Error, Unreachable, or Warning on those pages.


1. In the menu bar, go to the type of resource that has a status you want to acknowledge. For example, if you want to acknowledge the status of a switch, go to **Network > Switches**.



2. In the list of resources, locate the resource with the Error, Unreachable, or Warning status that you want to acknowledge.
3. Right-click the resource and select **Acknowledge status**. A window shows that the status for the specified resource was acknowledged.
4. Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge status**.
5. Click **Close**.

When you acknowledge the status of a top-level resource, the following actions are taken on pages in the web-based GUI:

- The dashboard:
  - The status of the acknowledged resource is not used to determine the status icon that is displayed for that resource on the resource diagram. For example, if you acknowledge the Unreachable status for a switch, that Unreachable status is not used to determine the status icon that is displayed for the switch icon on the dashboard.
  - The acknowledged status is shown as **Acknowledged** when you point to the status icon for a resource. For example, if 10 switches have an Unreachable status and you acknowledge the status for one of them, the following information is shown when you hover the mouse over the switch icon:

 **9 Unreachable (1 Acknowledged)**

**Restriction:** The dashboard shows status icons and hover information for servers, hypervisors, switches, fabrics, and storage systems. The statuses of pools and volumes are used to calculate the status of these top-level resources, so their individual statuses are not shown on the dashboard.

- The resource details page:
  - The status icons at the top of the page are updated. For example, if 10 switches have an Unreachable status, and you acknowledge one of them, the total number Unreachable statuses decreases by one and the number of Unreachable - Acknowledged statuses increases by one. The following information is updated at the top of the Switches page:

 **9 Unreachable**  **1 Unreachable - Acknowledged**

- The value in the **Status** column for the resource is updated to show an acknowledged status. For example, if you acknowledge the Unreachable status for a switch, the following value is shown in the **Status** column:

 **Unreachable**

**Related reference:**

“Resources that you can monitor in the web-based GUI” on page 128

You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

## Acknowledging the status of internal resources

The web-based GUI uses the status of internal resources to determine the status of top-level resources. Sometimes these statuses might represent problems that commonly occur or are known, but can be ignored. In such cases, you can acknowledge the status of an internal resource so that it is not used when determining the status of a top-level resource.

For example, if the status of a volume is Error, the status of the related storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall status of the storage system. In this case, if the other internal resources of the storage system are Normal, then the status of the storage system is also Normal.

1. In the menu bar, go to the type of resource in which the internal resource is located. For example, if you want to acknowledge the status of internal resources for a switch, go to **Network > Switches**.
2. Right-click a resource and select **View Details**. For example, if you want to acknowledge the status of a port on a switch, right-click that switch and select **View Details**.
3. In the **Internal Resources** section of the resource details page, click the name of the internal resource. For example, if you want to acknowledge the status of a port on a switch, click **Ports**.
4. Right-click the resource and select **Acknowledge status**. A window shows that the status for the specified resource was acknowledged. For example, if you want to acknowledge the status of a port, right-click the port and select **Acknowledge status**.
5. Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge status**.
6. Click **Close**. The status of the related, top-level resource might take several minutes to update.

When you acknowledge the status of an internal resource, the following actions occur:

- The status of the internal resource is no longer used to determine the status of the related, top-level resource.

For example, if you acknowledge the Error status for a controller on a server, the status of that controller is not used when determining the overall status of the server.

- The status icons at the top of the page for the internal resource are updated.

For example, if you acknowledge the Error status for a controller on a server, the total number of Error statuses decreases by one and the number of Error - Acknowledged statuses increases by one on the Controllers page. If originally there were five controllers with Error statuses, the following updated status information is shown at the top of the page:

 4 Error  1 Error - Acknowledged

**Related reference:**

“Resources that you can monitor in the web-based GUI” on page 128  
You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

---

## Monitoring the performance of resources


Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

In the web-based GUI, you can use the performance monitoring to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance of a resource fell outside of a specific range.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, Tivoli Storage Productivity Center correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including performance information.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

### Before you begin

Before you can view performance information for resources, you must complete the following tasks:

Task	GUI	 Learn more
Add storage systems and switches for monitoring by Tivoli Storage Productivity Center.	Web-based GUI	For information about how to add these resources, see the following topics: <ul style="list-style-type: none"> <li>• “Adding storage systems” on page 132</li> <li>• “Adding fabrics and switches” on page 133</li> </ul>
Ensure that performance data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	Web-based GUI	For information about verifying data collection and scheduling performance monitors, see the following topics: <ul style="list-style-type: none"> <li>• “Verifying that a performance monitor is running for a resource” on page 153</li> <li>• “Creating performance monitors” on page 152</li> </ul>
Define performance thresholds to monitor if the performance of a resource falls outside a specified range and might represent a potential problem.  When you define a performance threshold for a resource, select a specific metric that you want to measure. For example, you can define a threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.	Web-based GUI	For information about how to define performance thresholds, see the following topics: <ul style="list-style-type: none"> <li>• “Defining a performance threshold alert for a resource” on page 164</li> <li>• “Defining a performance threshold alert for multiple resources” on page 166</li> </ul>

**Related reference:**

“Performance thresholds” on page 1100

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

“Performance metrics” on page 1069


Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.


## Viewing performance information

View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

Before you view the performance of resources, ensure that performance data was collected for those resources during the time ranges that you want to analyze. Tivoli Storage Productivity Center uses performance monitors to collect metrics for measuring the performance of storage systems and switches. For information about how to verify that performance monitors were run, see “Verifying that a performance monitor is running for a resource” on page 153.

You can access the performance view for a resource from different locations in the web-based GUI:

Task	Accessing the performance view in the web-based GUI
<p>View the performance of specific storage systems, switches, and their internal resources.</p>	<ol style="list-style-type: none"> <li>1. In the menu bar, go to the type of resource that you want to view. For storage systems, go to <b>Storage &gt; Storage Systems</b>. For switches, go to <b>Network &gt; Switches</b>.</li> <li>2. Right-click a resource and select <b>View Performance</b>. By default, the following information is displayed in a chart for each type of resource: <ul style="list-style-type: none"> <li>• For storage systems, the top 5 storage systems with the highest total I/O rate are shown.</li> <li>• For switches, the top 5 switches with the highest total port data rate are shown.</li> </ul> </li> <li>3. To view the performance of an internal resource, right-click the top-level resource in the chart legend and select the option for the internal resource that you want to view. For example, for a SAN Volume Controller, right-click its name in the chart legend and select <b>MDisk Performance</b> to view the performance of its managed disks. For managed disks, the top 5 contributors to the storage system workload are shown on the chart. The default metrics are total back-end I/O rate and total back-end response time.</li> </ol> <p><b>Tips:</b></p> <ul style="list-style-type: none"> <li>• To add or change the metrics that are displayed for a resource, click the select metrics icon  and select the metrics to display.</li> <li>• You can also view the performance of internal resources on the details page for the related, top-level resource. For more information, see Table 59 on page 209.</li> <li>• To view the performance of multiple resources at the same time, press Shift or Ctrl and click the resources that you want to view. Press Shift and click to select consecutive rows in the list of resources; press Ctrl and click to select non-consecutive rows. Then, right-click any of the selected resources and select <b>View Performance</b>. Each resource is represented by a separate line in the performance chart and separate rows in the table view.</li> </ul>

Task	Accessing the performance view in the web-based GUI
View the overall performance of resources of the same type.	<ol style="list-style-type: none"> <li>In the menu bar, go to the type of resource that you want to view. For storage systems, go to <b>Storage &gt; Storage Systems</b>. For switches, go to <b>Network &gt; Switches</b>.</li> <li>Click the <b>Performance</b> tab. By default, the following information is displayed in a chart: <ul style="list-style-type: none"> <li>For storage systems, the top 5 storage systems with the highest total I/O rate are shown.</li> <li>For switches, the top 5 switches with the highest total port data rate are shown.</li> </ul> </li> </ol> <p><b>Tip:</b> To add or change the metrics that are displayed for a resource, click the select metrics icon  and select the metrics to display.</p>
View the real-time performance of a SAN Volume Controller or Storwize storage system.	<ol style="list-style-type: none"> <li>In the menu bar in the web-based GUI, go to <b>Storage &gt; Storage Systems</b>.</li> <li>Right-click a storage system and select <b>View Real-Time Performance</b>. A graphical representation of the processor usage of the system, and the bandwidth of the volumes, interfaces, and MDisks of the system is displayed.</li> </ol> <p>You can view the real-time performance of a storage system when the following conditions apply:</p> <ul style="list-style-type: none"> <li>The storage system is a SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V5000, Storwize V7000, or Storwize V7000 Unified system.</li> <li>The storage system must be at microcode level 7.3 or later.</li> <li>The storage system is running. To test the connection, right-click the storage system, and select <b>Connections &gt; Test Connection</b>.</li> <li>The storage system must be available to the system where the browser is running, through the local area network. To check if the storage system is available, right-click the storage system and click <b>Open Storage System GUI</b>.</li> </ul>
View the threshold violations that were detected during data collection and the performance of the resources where the violation occurred. You can use this method to view all threshold violations, not just the violations that generate alerts.	For information about how to view threshold violations, see “Viewing threshold violations” on page 219.

Task	Accessing the performance view in the web-based GUI
View the alerts that were generated by threshold violations and the performance of the resources where the violation occurred. Use this method to only view the alerts that were generated by threshold violations.	For information about how to view the alerts that are generated by threshold violations, see “Viewing alerts for threshold violations” on page 220.
View the performance of a resource that is included in a specific performance monitor.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, go to <b>Home &gt; Performance Monitors</b>.</li> <li>2. In the Name column, locate the name of the resource that you want to view.</li> <li>3. Right-click the performance monitor and select <b>View Performance</b>.</li> </ol>
View predefined reports and create custom reports that show the performance of resources. In Tivoli Common Reporting, you can view reports about storage systems, components of storage systems, switches, and switch ports.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, go to <b>Cognos &gt; View predefined reports</b>.</li> <li>2. To view predefined reports, navigate to the performance report that you want to view and click the name of the report.  To create a custom report about the performance of resources, in the menu bar, go to <b>Cognos &gt; Create performance reports</b>.</li> </ol>

In the web-based GUI, you can also view the performance of internal resources on the details pages for the related, top-level resources. For example, you can view the performance of the controllers for a storage system on the details page for that storage system. To access the performance view for an internal resource on a details page, select the type of internal resource and click the **Performance** tab.

The following information summarizes the internal resources and locations where you can access a performance view.


*Table 59. Internal resources and locations where you access the performance view*


Internal resource	Locations in the GUI where you can access the performance view
Pools	<ul style="list-style-type: none"> <li>• <b>Storage &gt; Pools</b></li> <li>• <b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Pools</b></li> </ul> <p>Available only for pools that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, and XIV systems.</p>
Volumes	<ul style="list-style-type: none"> <li>• <b>Storage &gt; Volumes</b></li> <li>• <b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Volumes</b></li> </ul>
Controllers	<b>Storage &gt; Storage Systems</b> , right-click the storage system, click <b>View Details</b> , and click <b>Controllers</b>

Table 59. Internal resources and locations where you access the performance view (continued)

Internal resource	Locations in the GUI where you can access the performance view
Disks	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Disks</b></p> <p>Available only for disks that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.</p>
Host Connections	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Host Connections</b></p> <p>Available only for host connections that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, and XIV systems.</p>
I/O Group	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>I/O Groups</b></p>
Modules	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Modules</b></p>
Nodes	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Nodes</b></p>
Managed Disks	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Managed Disks</b></p> <p>Available only for nodes that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.</p>
Ports on storage systems	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>Ports</b></p>
Ports on switches	<p><b>Network &gt; Switches</b>, and click <b>Ports</b></p>
RAID Arrays	<p><b>Storage &gt; Storage Systems</b>, right-click the storage system, click <b>View Details</b>, and click <b>RAID Arrays</b></p> <p>Available only for RAID arrays that are associated with System Storage DS series and ESS storage systems.</p>

**Related information:**

 Performance statistics (for SAN Volume Controller systems)  
 Real-time performance statistics provide short-term status information for the SAN Volume Controller system. The statistics are shown as graphs in the management GUI.

 Performance statistics (for Storwize V7000 Unified systems)  
 Real-time performance statistics provide short-term status information for the Storwize V7000 Unified system. The statistics are shown as graphs in the management GUI.

**How performance information is displayed**

The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.



## Performance chart

The top section of the performance view shows information about the selected resources. You can view this information in the following formats:

### Chart

The chart shows a visual representation of how the performance of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

### Table

The table shows performance and asset information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

## Performance chart legend

The bottom section of the performance view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides asset and performance information about a resource.

Each metric on the chart is represented by two columns in the legend. The first column includes the weighted average of the metric for each resource in the selected time range. The second column represents the peak data point of the metric for each resource in the selected time range. For example, if a metric in the chart is Total I/O Rate, the following columns are displayed in the legend: Avg Total I/O Rate, Worst Total I/O Rate. The worst value can be the minimum or the maximum number, depending on the metric. For example, the worst value for an I/O rate metric is a maximum number, while the worst value for the Cache Hit Percentage metric is a minimum number.

You can also view information about the resources that are related to the resources in the legend. For example, if a SAN Volume Controller is listed in the legend, right-click it to view the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage. For more information about the resources that you can view in the chart legend, see “Resources in the chart legend” on page 216.

### Tip:

- If you hide a resource in the performance chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this

information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

For example, if the current date is January 1, and you view a volume with the time range set to December 1 to December 7, the legend shows the capacity of that volume on December 7. If the capacity of the volume was changed between December 7 and January 1, this historical capacity is different from the current capacity of the volume on January 1.

#### **Related concepts:**

“Customizing lists of resources, tasks, and alerts” on page 287

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

#### **Controls:**

Each performance view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



#### **View chart**

View performance information in a chart format. The chart shows a visual representation of how resource performance trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of performance information at a specific time.



#### **View table**

View performance information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that is available.

For information how to filter, sort, and customize the columns in a table, see “Customizing lists of resources, tasks, and alerts” on page 287.



#### **Metrics**

Add metrics to the performance chart. On the Select Chart Metrics window, metrics are organized into three tabs: Volume Metrics, Disk Metrics, and Port Metrics. The metrics and tabs that are available depend on the type of resource that is being shown in the chart. The number next to the name of the tab represents the number of metrics that are currently selected from that tab.

You can select multiple metrics at the same time and from different tabs, but you cannot include more than two unit types in the same view. For

example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see “Performance metrics” on page 1069.




### Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.



### Specify granularity

Determine the granularity of the data that is shown in a performance view. Granularity determines the points that are shown on the y-axis of a chart, the points on the lines in a chart, and the rows in a table. When you first view performance information, the default time range is the last 12 hours and the granularity matches the frequency of data collection by the performance monitor (sample).

To change the granularity of chart or table, select one of the following options from the granularity menu:



#### (daily)

Set the granularity to show 1-day increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 day.



#### (hourly)

Set the granularity to show one-hour increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 hour.



#### (Sample)

Set the granularity to match the frequency of data collection by a performance monitor. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments.

This granularity is only available if the time range of the chart is less than 2 weeks.

If the performance monitor collects sample data at 1-minute intervals, data is displayed as 1-minute intervals in the chart only when data is available for all of the selected time range. If data is not available at 1-minute intervals for any portion of the selected time range, the data is displayed in 5-minute intervals.

For example, the performance monitor for a resource collects sample data every 5 minutes. Six hours ago, the performance monitor was changed to start collecting data every minute. You want to view sample data for the resource for the last 12 hours. However, the performance monitor did not collect sample data every minute for 6 of the last 12 hours. Therefore, the data is displayed in 5-minute intervals rather than 1-minute intervals.

**Restriction:** Each granularity has a maximum time range for displaying data. This time range is dynamically calculated based on the resources and metrics that are shown on the chart, and the data points per line. When the maximum range is exceeded, the granularity icon is disabled and a message indicates that there are too many lines on the chart. To avoid this problem, reduce the time range or the number of resources or metrics that are selected.

### Specify a time range

The time range of the performance information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the performance view, the default time range is the last 12 hours.

Ensure that performance data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



### Hide and show chart controls

The performance view lists the names of resources that are being shown in the chart and the controls for customizing that chart. Click the left arrow icon to hide the controls and display the chart in the entire window. Click the right arrow icon again to show the controls.



### Export information about the chart to a file

Export information on a performance view to a CSV file.



### Open the performance view in a separate web browser window

Open a duplicate of the current performance view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original performance view for comparison.



### **Synchronize the time range across all the open performance views**

Synchronize the time range across all the performance views that are displayed in separate browser windows. Use this action when you change the time range in a performance view and want to apply the same time range to the other performance views.

**Restriction:** This action affects only the performance views that are displayed in separate browser windows. The time range of the performance view that is shown as part of the main window for the GUI is not affected.

For example, you can view the storage system volumes that are assigned to a server. You can then open separate browser windows for the host connections, pools, and managed disks that are related to one of the volumes. If you change the time range in the performance view for host connections, click the synchronize button to apply the same time range to the views for pools and managed disks. The performance view of the original volume is not changed.

### **Actions for managing the resources in the chart legend**

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the **Actions** menu:

#### **View Properties**

View key details about a resource, including asset, status, configuration, capacity, and performance information.

#### **View *resource* Performance**

View the performance of resources that are internal or related to a resource in the chart legend. For example, when you view the performance of a SAN Volume Controller, you can right-click it and view the performance of its internal resources, such as disks, volumes, ports, managed disks, and back-end storage.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the performance view in the main window of the GUI.

**Restriction:** If performance data was not collected for an internal or related resource, that resource is not shown in the view. For example, if you right-click a storage system and select **Volume performance**, only the volumes for which performance data was collected are shown.

For more information about the resources that you can view in the chart legend, see “Resources in the chart legend” on page 216.

#### **Key Metrics View**

View a set of predefined metrics for resources in the chart legend.

#### **Custom View**

Customize the set of metrics that are shown for resources in the chart legend.

## Resources in the chart legend:

In the performance view, performance metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, you can open separate performance views for their internal or related resources. The separate performance views use the same time range as the performance view in the main window of the GUI.

For example, when you view the performance of a SAN Volume Controller, you can right-click it in the chart legend to open a performance view for its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.

The type of resource determines the internal and related resources that you can view. The following tables show the internal and related resources that are available when you right-click a resource in the chart legend.

*Table 60. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend*

<b>Resource in the chart legend</b>	<b>Right-click the resource in the chart legend to view the performance of the following resources</b>	<b>Right-click the resource in the chart legend to view the performance of the following related resources</b>
Storage system	Volume, RAID array, pool, controller, host connection, port	Storage virtualizer, switch
Volume	Storage system, RAID array, pool, controller, host connection	Managed disk
RAID array	Storage system, volume, pool, controller, host connection	None
Pool	Storage system, volume, RAID array, controller, host connection	None
Controller	Storage system, volume, RAID array, pool, host connection	None
Host connection	Storage system, volume, RAID array, pool, controller	None
Storage system port	Storage system	Switch port

*Table 61. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend*

<b>Resource in the chart legend</b>	<b>Right-click the resource in the chart legend to view the performance of the following resources</b>	<b>Right-click the resource in the chart legend to view the performance of the following related resources</b>
Storage system	Volume, pool, node, I/O group, disk, managed disk, host connection, port	Back-end storage system, storage virtualizer, switch

Table 61. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend (continued)

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Volume	Storage system, pool, node, I/O group, disk, managed disk, host connection	None
Pool	Storage system, volume, node, I/O group, disk, managed disk, host connection	None
Node	Storage system, volume, pool, I/O group, disk, managed disk, host connection, port	None
I/O group	Storage system, volume, pool, node, disk, managed disk, host connection, port	None
Disk	Storage system, volume, pool, I/O group, node, managed disk, host connection	None
Managed disk	Storage system, volume, pool, node, I/O group, disk, host connection	Back-end volume
Host connection	Storage system, volume, pool, node, I/O group, disk, host connection	None
Storage system port	Storage system, I/O group, node	Switch port

Table 62. Resources that you can view for XIV systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk
Host connection	Storage system, module, pool, volume	None
Storage system port	Storage system, module	Switch port

Table 63. Resources that you can view for SMI-S BSP systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, port	Storage virtualizer, switch
Volume	Storage system	Managed disk
Storage system port	Storage system	Switch port

Table 64. Resources that you can view for switches in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Switch	Port	Storage system, switch
Port	Switch	Storage system port, switch port

Table 65. Resources that you can view for servers in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Server	None	Storage system, volume, pool, controller, module, node, I/O group, disk, managed disk, host connection, switch, storage system port, switch port

## Saving URLs for performance views

You can create a bookmark or a favorite for a performance view.

The link in the bookmark or favorite stores details of the performance view, such as the resources, metrics, interval, and other details. The link also stores the time period that is specified in the time selector at the top of the chart.

You can also copy the URL from the address bar if you want to share the URL with colleagues.

Do one of the following steps:

- In Firefox, right-click in the performance view and click **Bookmark This Page**.
- In Internet Explorer, right-click in the performance view and click **Add to favorites**.

## Viewing threshold violations and alerts

View the performance thresholds that were violated and the alerts that were generated as a result of violations.

To view threshold violations and alerts for a resource, you must collect performance data and define performance thresholds for that resource. For



information about how to collect performance data, see “Creating performance monitors” on page 152. For information about how to define performance thresholds, see “Defining a performance threshold alert for a resource” on page 164.

A threshold violation is detected when a performance monitor discovers that the measured value of a metric falls outside of the specified threshold range for a resource. In the web-based GUI, you can view a list of all the threshold violations that are detected during data collection. You can also view all of the alerts that are generated for threshold violations.

## Viewing threshold violations

View the threshold violations that were detected on monitored resources during data collection. You can also open a threshold violation to view information about the resource where the violation occurred.

1. Complete one of the following tasks to view threshold violations at different locations in the web-based GUI:

Task	Steps
View all threshold violations	<ol style="list-style-type: none"> <li>1. In the menu bar, go to <b>Home &gt; Performance Monitors</b>.</li> <li>2. Click the <b>Threshold Violations</b> tab.</li> </ol>
View threshold violations for a specific type of resource	<ol style="list-style-type: none"> <li>1. In the menu bar, go to the type of resource. For example, if you want to view the threshold violations that were detected on storage systems, go to <b>Storage &gt; Storage Systems</b>.</li> <li>2. Click the <b>Threshold Violations</b> tab.</li> </ol>
View threshold violations for a specific resource	<ol style="list-style-type: none"> <li>1. In the menu bar, go to the type of resource. For example, if you want to view the threshold violations that were detected on a storage system, go to <b>Storage &gt; Storage Systems</b>.</li> <li>2. In the list of resources, right-click the resource whose threshold violations you want to view and select <b>View Details</b>.</li> <li>3. Click <b>Threshold Violations</b> in the <b>General</b> section.</li> </ol>


2. Right-click the threshold violation that you want to view and select **View Threshold Violation**. In this view, a chart shows the performance of the resource that violated the threshold. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:

- The blue line represents the performance of the resource for the selected metric.
- The yellow line represents the warning boundary value for the threshold.
- The red line represents the critical boundary value for the threshold.
- The dotted line represents the time when the violation occurred.

For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:

- A stress boundary was violated. In general, stress boundaries are upper-threshold boundaries that are violated when performance equals or exceeds a specified value.
- The threshold was violated on an internal resource other than a port.
- The threshold was violated on any of the following storage systems:
  - DS6000
  - DS8000
  - XIV system
  - ESS
  - SAN Volume Controller
  - Storwize V3500
  - Storwize V3700
  - Storwize V7000
  - Storwize V7000 Unified

You can also view the servers or hosts to which those volumes are assigned.

3. Optional: To access the full performance view of a resource that violated the threshold, click the open performance view icon  on the chart.

### Viewing alerts for threshold violations

View the alerts that were generated by threshold violations. You can also open an alert to view information about the resource where the threshold violation occurred.

**Tip:** Not all threshold violations generate alerts. A performance threshold must be defined to generate an alert before an alert can be displayed when the threshold is violated. For information about how to define a performance threshold with an alert, see “Defining a performance threshold alert for a resource” on page 164.


1. In the menu bar in the web-based GUI, go to **Home > Alerts**.
2. Right-click the threshold violation that you want to view and select **View Alert**. In this view, a chart shows the performance of the resource that violated the threshold and generated an alert. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:
  - The blue line represents the performance of the resource for the selected metric.
  - The yellow line represents the warning boundary value for the threshold.
  - The red line represents the critical boundary value for the threshold.
  - The dotted line represents the time when the violation occurred.

For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:

- A stress boundary was violated. In general, stress boundaries are upper-threshold boundaries that are violated when performance equals or exceeds a specified value.
- The threshold was violated on an internal resource other than a port.
- The threshold was violated on any of the following storage systems:
  - DS6000
  - DS8000

- XIV system
- ESS
- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

You can also view the servers or hosts to which those volumes are assigned.

3. Optional: To access the full performance view of a resource that generated an alert, click the open performance view icon  on the chart.

## Monitoring and administering applications and departments

You create applications to monitor the storage capacity, space usage, and performance of applications, and you create departments to monitor the space usage of the applications in the department. You can structure your applications and departments hierarchically to match the structure of your business organization.

You can map the application to the storage resources that the application uses to do capacity trending, health monitoring, and performance troubleshooting tasks. For example, a department might use 15 applications and be part of another five departments. Also a department might share storage resources with another department, subdepartment or application even if they do not belong in the same business hierarchy. Another example, might be if you are contacted because an application within a department is experiencing a performance issue. You carry out tasks like capacity reporting or trending on the behalf of the application.

### Application administration

Go to the pages for creating and administering the applications that you want to monitor.

*Table 66. Create and administer applications and subcomponents*

Actions	Navigation
Create applications and assign storage resources directly to the application.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. On the Applications page, click <b>Create Application</b>.</li> <li>3. On the Select a Method page, click <b>Assign storage resources to the application</b>.</li> </ol>
Create subcomponents that can have its own storage resources.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. On the Applications page, click <b>Create Application</b>.</li> <li>3. On the Select a Method page, click <b>Create a subcomponent that can have its own set of storage resources</b>.</li> </ol>

Table 66. Create and administer applications and subcomponents (continued)

Actions	Navigation
Create a filter to use in assigning resources.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. On the Applications page, right-click a row in the Applications list and click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Filters</b>.</li> <li>4. On the Filters page, click <b>Create Filter</b>.</li> </ol>
Add applications as subcomponents to an application.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. Right-click the applications that you want to add as subcomponents, and then click <b>Add to.. &gt; Application</b>.</li> <li>3. On the Add to Application page, select the applications and click <b>Save</b>.</li> </ol>
Add applications to departments.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. On the Applications page, select one or more applications, right-click and then click <b>Add to.. &gt; Department</b>.</li> <li>3. On the Add to Department page, select the applications that you want to add as members and click <b>Save</b>.</li> </ol>
Add applications to subdepartments.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. Right-click the department, and then click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Subdepartments</b>.</li> <li>4. Right-click the subdepartment, and click <b>View Details</b>.</li> <li>5. In the <b>General</b> section, click <b>Applications</b>.</li> <li>6. Click <b>Create Application</b>.</li> </ol>
Remove applications and subcomponents that you do not want to monitor.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. On the Applications page, select one or more applications, and click <b>Remove</b>.</li> <li>3. If you select <b>Remove Subcomponents</b>, the subcomponents that belong only to the application are removed.</li> </ol>
Remove subcomponents within the direct context of an application that you do not want to monitor.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. Right-click the application and click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Subcomponents</b>.</li> <li>4. Right-click the subcomponent, and click <b>Remove from Application</b>.</li> </ol>

Table 66. Create and administer applications and subcomponents (continued)

Actions	Navigation
Remove resources within the direct context of an application.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Applications</b>.</li> <li>2. Right-click the application and click <b>View Details</b>.</li> <li>3. In the <b>Related Resources</b> section, click an eligible resource.</li> <li>4. On the resource page, right-click the resource, and click <b>Remove from Application</b>.</li> </ol>

## Department administration

Go to the pages for creating and administering the departments that you want to monitor.

Table 67. Create and administer departments and subdepartments

Actions	Navigation
Create departments and associate members (applications and departments) to create a hierarchal structure.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. On the Departments page, click <b>Create Department</b>.</li> <li>3. On the Add Members page, click <b>Add one or more applications</b> or <b>Add one or more departments</b>.</li> </ol>
Add departments as subdepartments to a department.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. Right-click the departments that you want to add as subdepartments, and then click <b>Add to Department</b>.</li> </ol>
Add applications to departments.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. Right-click the department, and then click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Applications</b>.</li> <li>4. Click <b>Create Application</b>.</li> </ol>
Add applications to subdepartments.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. Right-click the department, and then click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Subdepartments</b>.</li> <li>4. Right-click the subdepartment, and click <b>View Details</b>.</li> <li>5. In the <b>General</b> section, click <b>Applications</b>.</li> <li>6. Click <b>Create Application</b>.</li> </ol>

Table 67. Create and administer departments and subdepartments (continued)

Actions	Navigation
Remove departments, subdepartments, and applications, that are members of the department.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. On the Departments page, right-click one or more departments and click <b>Remove</b>.</li> <li>3. If you select <b>Remove subdepartments and applications</b>, the subdepartments and applications that belong only to the department, that is being removed, are removed.</li> </ol>
Remove applications from a subdepartment.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. Right-click the department, and then click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Subdepartments</b>.</li> <li>4. Right-click the subdepartment, and click <b>View Details</b>.</li> <li>5. In the <b>General</b> section, click <b>Applications</b>.</li> <li>6. Select and right-click the applications, and then click <b>Remove</b>.</li> </ol>
Remove subdepartments within the direct context of a department that you do not want to monitor.	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, select <b>Groups &gt; Departments</b>.</li> <li>2. Right-click the department, and then click <b>View Details</b>.</li> <li>3. In the <b>General</b> section, click <b>Subdepartments</b>.</li> <li>4. Right-click the subdepartments that you want to remove, and then click <b>Remove from Department</b>.</li> </ol>

**Related tasks:**

“Monitoring capacity usage at different levels of a business hierarchy” on page 102  
 The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

“Comparing storage usage in each department” on page 104

In this scenario, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

“Using applications and subcomponents to monitor capacity and space usage” on page 106

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchal manner.

“Viewing storage capacity and usage trends” on page 108

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

## Applications

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

An application can be part of a department and have subcomponents (member applications) that can contain a five level deep hierarchy. Applications that are grouped together can range from large line of business systems to specialized software, in a department that runs on either client computers or servers. For example, an application might be an automated billing system in the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

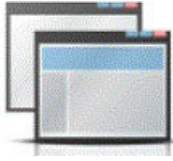
Another example might be an application named *Database* and the database contains logs and data, and a best practice is to place the logs and data on separate volumes. You might create subcomponents for the logs and data and these subcomponents become part of the *Database* application.

A storage administrator can use applications for planning purposes to diagnose how much storage a particular application is consuming in relation to the other applications in the business hierarchy. Another applicable usage for applications might be troubleshooting performance issues. The administrator checks the related resources that make up the application to see if there is a health status issue or can examine the performance of volumes that are used by the application, and, then troubleshoot along a particular data path.

The Applications page shows capacity information about the applications that are monitored by Tivoli Storage Productivity Center and the actions that you can use to view and manage them. To see more detailed information, view properties and details for the application. If the application is associated with a department, information about that department is also shown.

The following information shows the resources and objects that can be added to the application.

Table 68. Resources and objects that can be added to the application

Object	Resources or objects that can be added to the application
 <p data-bbox="467 1612 592 1644">Application</p>	<ul style="list-style-type: none"> <li data-bbox="829 1438 992 1465">• Applications</li> <li data-bbox="829 1474 948 1501">• Volumes</li> <li data-bbox="829 1509 987 1537">• Hypervisors</li> <li data-bbox="829 1545 976 1572">• Data Stores</li> <li data-bbox="829 1581 932 1608">• Servers</li> <li data-bbox="829 1617 922 1644">• Shares</li> <li data-bbox="829 1652 1382 1717">• File Systems (Only server file systems monitored through a Storage Resource Agent.)</li> <li data-bbox="829 1726 1024 1753">• Volume Groups</li> </ul>

### Creating applications

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

Use the web-based GUI to create applications, assign resources and subcomponents (member applications) for monitoring by Tivoli Storage Productivity Center.

To determine the storage capacity usage by the storage resources that the applications use, you create a data model of the applications for your storage environment.

To create an application, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups > Applications**.
2. On the Applications page, click **Create Application**.
3. On the Create Application page, enter information about the application including name, description, type, and subtype.

**Note:** You cannot have a duplicate application name on the same level in your application hierarchy. For example, you create an application named *Finance* and create a subcomponent named *Business*. You can create another application named *Business* as a top level application, but you cannot add the new *Business* application as a subcomponent of the *Finance* application.

You can customize the icon for your application by clicking the existing icon on the Create Application page and selecting another icon from the Customize Icon page.

4. Click **Create**.
5. Optional: On the Select a Method page, select **Assign storage resources to the application that was created** and follow the instructions to add the resources directly to the application.
6. Click **Preview** to see a list of filter matches before you finalize your filtering criteria, otherwise, click **Save**.
7. Optional: Select **Create a subcomponent that can have its own set of storage resources** and follow the instructions.
8. Click **Create**.

When the application creation is created, you can view status information, block and file capacity and additional data on the Applications page and in properties. For a detailed view of the application and related resources, on the Applications page, right-click the application and select **View Details**. To see the subcomponents for the application, click **Subcomponent** on the application details page.

**Related tasks:**

“Adding subcomponents” on page 227

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

“Using filters to assign resources to applications” on page 227

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

“Viewing information about applications” on page 230

You can view information about the application that you created such as related resources, filters that are used to populate the application with resources, any subcomponents and capacity trending and performance information.



## Adding subcomponents

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

To create a subcomponent (member application) for an existing application, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups > Applications**.
2. On the Applications page, select one or more applications, right-click and then, click **Add to.. > Application**.
3. On the Add to Application page, choose the applications that you want to become subcomponents of your selected applications.
4. Click **Save**.
5. Optional: On the details page for the application, click **Subcomponents**.
6. On the Subcomponents page, click **Create Subcomponent**.
7. On the Create Subcomponent page, enter information including name, description, type, and subtype.
8. Click **Create**.

After the subcomponent is created, you can view the results on the Subcomponent page and in properties. For a more detailed view of the subcomponent, select the subcomponent from the Subcomponents page and click **Actions > View Details** .

## Adding applications to departments

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

To add applications to departments, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups > Applications**.
2. On the Applications page, select one or more applications, right-click and then, click **Add to.. > Department**.
3. On the Add to Department page, choose the applications that you want to add as members to the department.
4. Click **Save**.

You can view the added departments on the **Departments** page.

### Related tasks:

*“Creating applications” on page 225*

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

*“Creating departments” on page 233*

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in Tivoli Storage Productivity Center.

## Using filters to assign resources to applications

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

When you create a filter to assign storage resources to applications, the resources that match the criteria that are defined in the filter are automatically added to the application or subcomponent (member application).

The use of wildcard characters enables you to customize and expand your search to add related resources. The asterisk (\*) and the question mark (?) can be used as wildcard characters. The asterisk matches one or more characters. The question mark matches a single character. Use a comma-separated list to add multiple resources. As new resources are added or modified and they match the existing filter, they automatically become part of the application. For example, a new server is added, a volume is renamed or a share is mounted on a server. If you edit a filter, the resources that no longer match the filter are removed from the application and the resources that now match the filter are added to the application. If you remove the filter, all storage resources that match the filter criteria are removed from the application.

Filtering is extremely useful in storage environments where resources are predefined and named in a consistent format. Instead of adding each resource separately, you create filtering criteria that matches your naming format in your environment. You can use filtering when you create a new application or directly from the details page for an existing application.

You can only select one resource type filter at a time during the creation workflow, but you can create multiple filters for applications.

The filters are divided into the following categories:

#### **All storage that belongs to a set of servers/hypervisors**

You can add servers or hypervisors to the application or subcomponent (member application).

For example, in the **Define Name Patterns** section, enter \*blade, in the **Server/Hypervisor** field and click **Save**. The filter adds all known servers or hypervisors that have a name that ends with blade to the application.

If you enter server???.abc.com, in the **Server/Hypervisor** field, the filter adds all servers, in the domain with names that begin with server followed by three characters and ends with .abc.com to the application. If you enter, accounting\*.\*, the filter adds all servers with names that begin with accounting regardless of the domain the system belongs to.

If you enter cs-blade47.lab.abc.com, cs-blade49.lab.abc.com or cs-blade47\*, cs-blade49\*, the filter adds multiple servers to the application.

#### **A selection of volumes that match a pattern**

You can add volumes of a certain name format, volumes that belong to a server or hypervisor, volumes from a specific storage system or volumes from a pool to the application or subcomponent. You can only select one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter \*vol, in the **Volumes** field, mark the **Belonging to** box, click **Server/Hypervisor**, enter \*.lab.abc.com and click **Save**. The filter adds all volumes that have a name that ends with vol and are mapped to servers that have a name that ends with .lab.abc.com to the application.

If you enter \*vol, in the **Volumes** field, mark the **Belonging to** box, click **Storage System**, and enter ds8000. The filter adds all volumes that have a name that ends with vol that belong to the ds8000 storage system.

#### **A selection of shares that match a pattern**

You can add shares and exports that are mounted on a server or hypervisor, or shares and exports that are exported from a specific storage system to the application or subcomponent. You can select only one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter \*share in the **Shares** field, mark the **Belonging to** box, click **Server/Hypervisor**, enter \*.lab.abc.com and click **Save**. The filter adds shares that have a name that ends with share that are mounted on servers that have a name that ends with .lab.abc.com .

If you enter share3 in the **Shares** field, mark the **Belonging to** box, click **Storage System** and enter SONAS, the filter adds all the shares that are named share3 and are exported from the SONAS storage system.

To create a filter, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups>Applications**.
2. On the Applications page, right-click a row in the Applications list and click **View Details**.
3. In the **General** section, click **Filters**.
4. On the Filters page, click **Create Filter**.
5. On the Create Filter page, select a resource type filter and enter the necessary information.
6. Click **Preview** to see a list of filter matches before you finalize your filtering criteria.
7. Click **Save** .

The **General** section of the details page shows the number of filters that you created, for example, **Filters (6)**. Click **Filters** to access the Filters page and see the filters that were created. For example, The number of matches: 26, which indicates there are 26 matches for the filter criteria. Click the link to see the resources that are available for the match.

#### **Related tasks:**

*“Creating applications” on page 225*

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies that the applications use, create applications for monitoring in Tivoli Storage Productivity Center.

*“Adding subcomponents” on page 227*

You created a new application and a subcomponent (member application) with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

### **Adding resources manually to applications**

You can add resources directly from a resource page, for example, servers and volumes, to your applications.

Only the following resources can be added directly to applications:

- Servers
- Hypervisors

- Volumes
- Volume groups
- Data stores
- Shares
- File systems

To add a resource directly to an application, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups > Servers**.
2. On the Servers page, select a server from the Servers list, right-click and click **Add to Application**.
3. On the Add to Application page, select one or more applications to assign the server to.
4. Click **Save**.

The resource is added to the application and is displayed on the Applications page.

**Related tasks:**

“Using filters to assign resources to applications” on page 227

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

**Adding resources with the command line interface**

You can use the **modifyappgroupviafile** command-line interface (CLI) command to automate your bulk assignment of eligible storage resources to applications. To add a single resource to an application, use the **modifyappgroup** command.

**Related information:**

 [modifyappgroupviafile command](#)

 [modifyappgroup command](#)

**Viewing information about applications**


You can view information about the application that you created such as related resources, filters that are used to populate the application with resources, any subcomponents and capacity trending and performance information.

To view the details of an application, use these steps:

1. In the menu bar in the web-based GUI, go to **Groups > Applications**.
2. On the Applications page, right-click an application and click **View Details**.

The application details shows information about the specific application. The information is available from the following sections on the page:

- **Overview:** The overall capacity trending about the application and performance information about volumes and switch ports.
- **Properties:** The key information about the application.
- **Filters:** The number of filters created and the filtering criteria used for the resources that were added to the application.
- **Subcomponents:** The number of subcomponents (member applications) that are a members of the selected application. If there is more than one subcomponent, the number of subcomponents is displayed.
- **Related Resources:** The information about the resources that are related to the application.

**Tip:** To view descriptions of the information that is available on an application details page, click the **Help** icon  in the upper-right corner of the page.

## Removing applications and subcomponents

You can remove applications and their subcomponents. The data that was collected to monitor the applications is also removed from the database repository.

To remove applications and subcomponents from a level of the business hierarchy, complete the following steps:

1. In the menu bar in the web-based GUI, select **Groups > Applications**.
2. On the Applications page, select one or more applications, and click **Remove**.
  - a. If you select **Remove Subcomponents**, the subcomponents that belong only to the application are removed.
  - b. If you do not select **Remove Subcomponents**, any subcomponents of the application, that is being removed, then become members of the parent application or department. If the application is a top level application, and, the subcomponents are not members of any other application, the subcomponents become top level applications.

The applications and subcomponents are removed and are no longer displayed on the Applications page.

To remove subcomponents from within the direct context of the application, use these steps:

1. In the menu bar in the web-based GUI, select **Groups > Applications**.
2. On the Applications page, select an application from the Application list and right-click **View Details**.
3. Optional: Select a application from the Applications page and right-click **View Properties**.
  - a. Click the **Subcomponents** tab and continue with steps 5-6.
4. On the application details page, click **Subcomponents** in the **General** section.
5. On the Subcomponents page, select the subcomponent and right-click **Remove from Application**.
6. Confirm that you want to remove the subcomponent and click **Remove**.

## Removing resources from applications

Related resources within the context of an application can be removed from applications. The action is only available when viewing the related resources that belong to an application on the details page.

The following resources can be removed from an application:

- Volumes
- Shares
- Servers
- File systems (only server file systems monitored through a Storage Resource Agent)
- Volume Groups
- Hypervisors
- Data Stores

Resources that were added through a filter match are removed if the filter is modified and no longer includes the resource that was removed from the application, the filter is removed from Tivoli Storage Productivity Center, or the resource is removed from Tivoli Storage Productivity Center. Resources can be added back to the application if the filter is modified and includes the resource that was originally removed, a new filter is created that matches the resource, the resources are added to Tivoli Storage Productivity Center, or you use the **Add to Application** action.

For example, a server named, abc.xyz.com is added to an application named, ApplicationABC because it matched the server name filter, abc\*. You decide to remove the abc.xyz.com server from the ApplicationABC application and remove the abc.xyz.com server from Tivoli Storage Productivity Center. Later you want to add the abc.xyz.com server back to Tivoli Storage Productivity Center. The abc.xyz.com server is added to the ApplicationABC application because the name matches the server resource filter that you originally created for the ApplicationABC application. If you do not want the resource added back to the application, you can further refine your filters. Click the **View Filters** link on the confirmation dialog to make any adjustments to the filtering criteria that you might need.

Resources that are automatically added to an application because they are related to resources that were explicitly added to the application *are not* removed from the application. The resources are automatically removed when the corresponding related resource that was added to the application, is also removed from the application.

For example, you add a server to the application that has one volume from a storage system. The volume is added as a related resource of that server in the application. If you attempt to remove the volume from the application, the volume will not be removed since it is a related resource. The volume is only removed when the server that was explicitly added to the application is removed.

To remove resources directly from an application, use these steps

1. In the menu bar in the web-based GUI, select **Groups > Applications**.
2. On the Applications page, select an application from the Application list and right-click **View Details**.
3. On the application details page, click an eligible resource in the **Related Resource** section.
4. On the resource page, select the resource from the list, right-click **Remove from Application**.
5. Confirm that you want to remove the resource.
6. Click **Remove**.

The resources that were added to the application are removed from the application.

**Related tasks:**

“Using filters to assign resources to applications” on page 227

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

## Departments


View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or application even if they do not belong in the same hierarchy. For example, in a collaboration scenario, a single IBM SAN Volume Controller might be shared by multiple departments. You can provide reports on the storage capacity that is used at different levels of your business hierarchy to regulate and compare the capacity storage, and in turn, track the trending of storage usage by specific departments.

The Departments page shows storage capacity information about the top level departments, the subdepartments and any applications that belong to the department that are monitored by Tivoli Storage Productivity Center and the actions that you can use to view and manage the resources that are defined to them. To see more detailed information, view properties and details for the department.

The following information shows the objects that you can add using the Add to Department function.

Table 69. Objects that can be added to the department

Object	Objects that can be added to the department
 Department	<ul style="list-style-type: none"><li>• Applications</li><li>• Departments</li></ul>

### Creating departments

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in Tivoli Storage Productivity Center.

Use the web-based GUI to create departments, subdepartments (member departments) and add applications for monitoring different levels of a business hierarchy.

To create a department, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups > Departments**.
2. On the Departments page, click **Create Department**.
3. On the Create Department panel, enter information about the department including name, description, type and subtype. You cannot have a duplicate department name on the same level in your department hierarchy. For example, you create a department named *Finance* and create a subdepartment named

*Business.* You can create another department named *Business* as a top level department, but if you want to make your new *Business* department a subdepartment of the *Finance* department, you see a message that indicates there is already a department named *Business* that is a member of the *Finance* department; your department creation is denied.

You can customize the icon for your department by clicking the existing icon on the Create Department page and selecting another icon from the Customize Icon page.

4. Click **Create**.
  - a. On the Add Members page, click **Add one or more applications** .
  - b. On the populate page, select the applications that to add to your department. If you do not want to add any of the applications that are listed, click **Create Application** to create a new application to add to the department.
  - c. Click **Finish**.
  - d. Click **Add one or more departments** .
  - e. On the populate page, select the departments that to add to your department. If you do not want to add any of the departments that are listed, click **Create Department** to create a new department to add to the department.
  - f. Click **Finish**.

To view the newly created department, go to the Departments page. To see the applications that you added to your department, see the details page and click **Applications**. If you added departments to a department, see the **Subdepartments** column on the Departments page.

**Related tasks:**

“Adding subdepartments”

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add subdepartments.

“Adding applications to departments” on page 227

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

## **Adding subdepartments**

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add subdepartments.

To add a subdepartment to the department level hierarchy, use these steps:

1. On the Departments page, select one or more departments, right-click and click **Add to Department**.
2. On the Add to Department page, select one or more departments from the list and click **Save**. The department or departments you select from the Add to Department page become parents of the selected department. You can also view the applications that are members of the department on the Departments page.
3. Optional: On the details page for the department, click **Subdepartments**.
4. On the Subdepartments page, click **Create Subdepartment**.
5. On the Create Subdepartment page, enter information including name, description, type, and subtype.
6. Click **Create**.



View the subdepartment on the Departments page or the details page.

**Related tasks:**

“Creating departments” on page 233

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in Tivoli Storage Productivity Center.

## Viewing information about departments


You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments (member departments) that are used to establish a business hierarchal structure.

To view the details of a department, use these steps:

1. In the menu bar in the web-based GUI, go to **Groups > Departments**.
2. On the Departments page, right-click a department in the list and click **View Details**.

The department details page shows information about the specific department. The information is available in the following sections on the page:

- **Overview:** The overall capacity trending and performance information about the department.
- **Properties:** The key information about the department.
- **Applications:** The applications that belong to the department. If there is more than one application, the number of applications is displayed.
- **Subdepartments:** The subdepartments that belong to the department. If there is more than one subdepartment, the number of subdepartments is displayed.
- **Related Resources:** Information about related resources in the storage environment that the applications, which belong to the department, are using or are connected to.

**Tip:** To view descriptions of the information that is available on a department details page, click the **Help** icon  in the upper-right corner of the page.

## Removing departments and subdepartments

You can remove departments, subdepartments (member departments) and applications that are members of the department. The data that is collected about the departments, subdepartments and applications is removed from the database repository.

To remove departments, subdepartments, and applications, that are members of the department, from a level of the business hierarchy, complete these steps:

1. In the menu bar in the web-based GUI, select **Groups > Departments**.
2. On the Departments page, select one or more departments, right-click and click **Remove**.
3. Confirm the information and click **Remove**.
  - a. If you select **Remove subdepartments and applications**, the subdepartments and applications that belong only to the department, that is being removed, are removed.
  - b. If you do not select **Remove subdepartments and applications**, the subdepartments and applications of the department that are being removed become members of the parent department. If the department is a top level

department, and the subdepartments are not members of any other department, they become top level departments.

The departments, subdepartments and applications are removed from the business level hierarchy and are no longer displayed on the Departments page.

To remove subdepartments from within the direct context of a department, use these steps:

1. In the menu bar in the web-based GUI, select **Groups > Departments**.
2. On the Departments page, select a department from the list and right-click **View Details**.
3. Optional: Select a department from the Departments page and right-click **View Properties**.
  - a. Click the **Subdepartments** tab and continue with steps 5-6.
4. On the department details page, double-click **Subdepartments** in the **General** section.
5. On the Subdepartments page, select the subdepartment and right-click **Remove from Department**.
6. Reply to the confirmation message and click **Remove**.

### **Removing applications from a department**

You can remove an application from within a department. This action is only available from the department detail page or the properties for the department.

To remove applications from within the context of a department, use these steps:

1. In the menu bar in the web-based GUI, select **Groups > Departments**.
2. On the Departments page, select a department from the list and right-click **View Details**.
3. Optional: Select a department from the Departments page and right-click **View Properties**.
  - a. Click the **Applications** tab and continue with steps 5-6.
4. On the department details page, click **Applications** from the **General** section.
5. On the Applications page, select the application and right-click **Remove from Department**.
6. Confirm the information and click **Remove**.

The applications that were added within the context of a department are eligible for removal from within department.

#### **Related tasks:**

“Viewing information about departments” on page 235

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments (member departments) that are used to establish a business hierarchal structure.

## **Exporting information about applications and departments**

You can save information about an application and department to a PDF, CSV, or HTML file. Information that you can export from the web-based GUI includes all the values that are being shown in the columns for a list of applications and departments.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the web-based GUI has a column or set of columns that determines the order of its rows.
- Filter the list to show only the applications and departments that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the following pages:

### **Applications and departments list and details pages**

You can export information about the top level applications and departments and the related resources that are shown on application and department list and detail pages.

For example, you can export information about monitored applications and their subcomponents and departments on the Applications page. When you select the export action, the generated file includes the column values for each monitored application in the list.

For example, you can export information about the servers that are associated with an application on the details page.

1. In the web-based GUI, go to the Applications or Departments page and select applications or departments for which you want to export information.
2. Click **Actions** and select **Export > file\_format**, where *file\_format* represents the format that you want to use for a file. The following formats are available:

#### **CSV (comma-separated values)**

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name","Column_name","Column_name","Column_name","Column_name","Column_name"
"data","data","data","data","data","data"
```

where `Column_name` represents the name of a column in a table and `data` represents the data that is associated with a column.

**PDF** You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

**Tip:** When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

## HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

### Related concepts:

“Customizing lists of resources, tasks, and alerts” on page 287

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

---

## Viewing the data paths of resources

A data path shows the connectivity between host connections, servers, hypervisors, virtual machines, storage systems, and the fabrics through which they share storage. You can use the data path view to monitor status and pinpoint problem areas in selected data paths. This view includes graphical and tabular representations of the top-level resources in a data path.

If a resource does not map storage to other resources, or the resource is not assigned storage from other resources, the data path view is not available.

To view the data path for a resource, complete these steps:

1. In the menu bar, select the type of resource that you want to view. For example, if you want to view the data path for a server, go to **Servers > Servers**.
2. Right-click the resource and select **View Data Path**.

Depending on how resources are configured, a data path can show between three to six types of top-level resources. For example, a data path with three types of top-level resources might show the following relationships:

server <> fabric <> storage system

A data path with six types of top-level resources might show the following relationships:

virtual machine <> hypervisor <> fabric <>  
storage virtualizer <> fabric <> storage system

Before you can view a data path, you must complete the following tasks:

- Add the resource that has the data path you want view. For information about how to add a resource, see “Adding resources” on page 127.

**Tip:** Resources that are not monitored by Tivoli Storage Productivity Center are also shown in a data path view. However, detailed information about their status and internal resources is not available.

- Schedule a probe to collect data about your monitored resources to ensure that the most recent information is shown for a data path. For information about probes, see “Collecting asset and status data” on page 147.

You can use the data path view to complete the following actions:

- View the path of data that is shared between resources. The data path view shows only the resources that have storage mapped between them. For example, if you view the data path for a storage system, only the host connections that are assigned volume space from that storage system are shown. Host connections that are connected to the storage system, but do not have storage assigned storage to them, are not shown.
- View the fabrics through which resources in a data path are communicating.
- View the propagated status of the top-level resources that are in a data path.
- View the status of the internal resources for top-level resources that are in a data path.
- Customize the appearance of the data path view to suit the needs of your environment.
- Export the data paths view as an image or CSV file.

Information about a data path is organized into two views: a graphical view and a table view. The graphical view shows icons for connected resources and uses lines to represent the path of data through them. The table view includes a row for each resource in a data path. The row is divided into columns that contain information about the resource.

## Resources that you can view in a data path

The **Topology View** tab in the data path view includes a visual representation of the data path for a top-level resource. In a data path, icons represent the connected resources and lines represent the path of data through them. For example, resources that are storage systems are shown with an icon that resembles a storage system.

The following table shows the icons for each type of resource that can be shown in a data path.

*Table 70. Resource icons in a data path*




Resource	Icon
Server (with agent)	
Server (without agent) <sup>1</sup>	
Hypervisor	

Table 70. Resource icons in a data path (continued)












Resource	Icon
Virtual machine	
Fabric	
Virtual fabric	
DS8000	
DS6000	
SAN Volume Controller	
Storwize V7000 <sup>2</sup>	
XIV system	

Table 70. Resource icons in a data path (continued)

Resource	Icon
ESS	
Other storage systems	
Unmanaged host connections	
<p><b>Tips and restrictions:</b></p> <ul style="list-style-type: none"> <li>• <sup>1</sup> When agentless servers and servers with agents are grouped together, the icon for servers with agents is used to represent that group.</li> <li>• <sup>2</sup> If a Storwize V7000 is not monitored by Tivoli Storage Productivity Center, but is connected to a monitored hypervisor, it is shown in the data path with the icon for a SAN Volume Controller. For example, when a virtual machine is monitored, a probe collects information about its disks. If the disks are attached to a SAN-attached storage system that is not monitored, Tivoli Storage Productivity Center uses the probe information to determine the type of storage system. If the storage system is a Storwize V7000, the icon for SAN Volume Controller is shown instead of the icon for a Storwize V7000.</li> <li>• Only the resources that are sharing storage in a data path are shown.</li> <li>• Servers and hypervisors that are not monitored by Tivoli Storage Productivity Center, but are assigned storage from a monitored storage system, are shown as unmanaged host connections. For example, if a storage system has a volume that is mapped to an unknown server, the server is represented by an unmanaged host connection in the view.</li> <li>• For storage systems with multiple host connections, the host connections are included in the same resource box if they communicate with the storage system through the same fabric. For example, if a storage system has volumes that are mapped to 20 host connections that use the same fabric, a single resource box represents the 20 host connections.</li> <li>• Storage systems can have volumes that are provisioned to more than one server in a clustered environment. Because all the clustered servers can be part of the data path, each server is shown in the view.</li> </ul>	

## The status of resources

Each resource in a data path is represented by a resource box. A resource box contains an icon that represents the overall status of a resource. This status is determined by the propagated status of the internal resources for the resource. For example, in a storage system, a port might have an Error status and a pool might

have a Warning status. In this case, the Error icon is shown on the resource box for the storage system because the Error condition is the most severe status that was detected on its internal resources.



Some resources in a data path might be shown with an unknown status. Resources can have an unknown status if they not monitored by Tivoli Storage Productivity Center, but are connected to other, monitored resources. Tivoli Storage Productivity Center uses the information that it collects about monitored resources to determine basic information about the connected resources that are not monitored.

For example, if a host connection is assigned volumes, but Tivoli Storage Productivity Center is not monitoring the fabric through with the host connection communicates with the storage system, the fabric is shown with an unknown status.



## Navigating in a data path

A data path might contain more resources that can be fit in the view at the same time. In these cases, you can scroll the view to show different parts of a data path.

1. In the **Topology View** tab, click the **Pointer** button .
2. Click an empty area in the view and drag the mouse pointer to scroll to different areas of a data path.
3. Optional: To view where you are in the overall topology of a data path, click the **Topology overview** button  in the upper-right corner of the topology view. The red box shows the visible area within the overall topology view, and status icons are displayed for the resources in the data path.

## Viewing information about top-level resources

You can view detailed information about the top-level resources in a data path. Top-level resources that can be shown in a data path include storage systems, servers, hypervisors, and fabrics.

1. To view key attributes of a specific top-level resource without leaving the **Topology View** tab, right-click the resource and select **View Properties**. A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
2. To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**. A resource details page is displayed and includes information about storage statistics, alerts, jobs, internal resources, and related resources.

## Viewing information about internal resources




You can view information about the internal resources that are associated with the top-level resources in a data path. Internal resources are components that exist in a top-level resource.



1. In the **Topology View** tab, hover the mouse pointer over a top-level resource in the data path.
2. View the following information about internal resources:
  - The number of internal resources that are associated with a top-level resource. For example, if a server has 20 disks, the value **Disks (20)** is shown.
  - The overall status of the internal resources. For information about how the overall status of internal resources is determined, see “How status is determined” on page 197.



## Zooming in

You can use the zoom buttons in the toolbar or the mouse to magnify areas in a data path.

1. To zoom in on a specific area in a data path, click the **Zoom to selected** button  and click and drag the mouse pointer over the area.
2. Optional: To use the mouse to zoom in on a specific area in a data path, move the mouse pointer over the area where you want to zoom in and scroll the middle mouse button forward.
3. Optional: To zoom in on the area that is in the center of the view, click the **Zoom in** button .
4. Optional: To adjust the zoom level of the view so that all the resources in a data path are shown at the same time, click the **Fit to view** button .

## Zooming out

You can use the zoom out button in the toolbar or the mouse to zoom out in the view.

1. To zoom out from the area that is in the center of the view, click the zoom out button .
2. Optional: To use the mouse to zoom out from a specific area in a data path, move the mouse pointer over the area where you want to zoom out and scroll the middle mouse button backward.
3. Optional: To adjust the zoom level of the view so that all the resources in a data path are shown at the same time, click the **Fit to view** button .

## Changing the layout

You can change the layout of the resources in a data path. The layout determines how resources are arranged in the view.

1. Click the **Layout** button in the toolbar. The icon that is displayed for the **Layout** button represents the current layout of the data path.
2. Select one of the following layouts:



### Hierarchical

Organize resources in a vertical format. Resources that use storage are shown at the top of the view. Resources that provide storage are shown at the bottom of the view.



### Tree

Organize resources in a vertical and horizontal layout.




### Customize

Customize the layout of a data path. When you customize a layout, use the hierarchical or tree layouts as the base and specify values for the associated options.

- Optional: If you selected **Customize**, click **OK** to apply your settings.


## Saving the display settings

You can save the display settings that are defined for a data path view. The next time that you view the data path for the resource, the same display settings are applied.

- In the **Topology View** tab, click the **Save** button .
- Optional: To use the keyboard to save the display settings, press Ctrl+S.
- Click **Save Customizations**.

## Printing a view

You can print the current view of a data path in the **Topology View** tab. The printing capability is not available in the **Table View** tab.

- Click the **Print** button .
- Specify the following options when printing:
  - In the **Page margin** field, enter the amount of space between the edges of the page and the view. The default setting is 1 inch.
  - In the **Maximum pages in width or height** field, specify the number of pages across which the content of a view is printed. This number determines the maximum number of pages that are printed. For example, if you enter **1**, the content of a view is printed on one page. If you enter **2**, the view is printed across four pages: two pages horizontally and two pages vertically.
- Click **OK**.
- Select the printer settings and click **OK** to print the view.

## Exporting a data path to a file

You can save the contents of a data path to an image file or CSV file for offline viewing.

- In the **Topology View** tab, right-click on an empty area in the data path.
- Complete one of the following actions to select an output format for the file:
  - Select **Export as Image** to save the data path as a JPEG or PNG file.
  - Select **Export** to save the data path to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information for a data path is organized into the following column in the CSV files: Name, Type, From, RelType, To.
- Specify the location and name of the file.
- Save the file.

---

## Provisioning storage

You can assign storage to servers, hypervisors, and clusters on servers and hypervisors with the Provisioning Storage wizard in the web-based GUI. You can assign storage to hypervisors with the vSphere Web Client extension for Tivoli Storage Productivity Center.

You can assign the following types of storage to servers, hypervisors, and clusters on servers and hypervisors:

- Volumes
- NAS file shares

When you assign storage, you must associate the storage with a service class. You can associate the storage with a default service class such as gold, silver, or bronze, or you can create a user-defined service class. You can also create capacity pools to restrict the provisioning of storage to a defined set of storage resources such as pools on a storage system in a particular location.

## Configuring Tivoli Storage Productivity Center for provisioning

Before you can use the simplified provisioning capabilities that are provided with Tivoli Storage Productivity Center, you must configure your environment.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, Tivoli Storage Productivity Center can analyze and optimize volume performance.

To configure Tivoli Storage Productivity Center for provisioning, complete the following tasks:

1. Assign storage pools to tiers by setting the tier level of each storage pool. The tier levels are number tags that can reflect any tier structure in the environment.

You must assign storage pools to tiers to later provision volumes that require a certain tier level and to optimize the placement of volumes. Storage pools do not need to be assigned to tiers to provision volumes that do not require a tier level, or to provision file shares.

2. Create service classes. A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.
3. Optional: Create capacity pools. You can restrict provisioning requests to a capacity pool. Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business, or to a particular application.

### Configuring service classes

Configure service classes to represent the different levels or types of storage quality that you want to provision. When you are configuring service classes for your installation, consider whether certain storage requests must always be satisfied from a particular set of resources. Also, consider whether you want to allow users without administrator privileges to provision storage, and whether you want their provisioning requests to require administrator approval.

## Creating service classes:

Create service classes to specify storage requirements for future provisioning requests. A service class describes attributes and capabilities of block storage or file storage resources, and typically describes particular quality of service.

You must have Administrator privileges to create a service class.

The Create Service Class wizard guides you through the steps to create a service class. When you provision a volume or share, you define storage requirements by specifying a service class. A block-storage service class specifies the requirements for the block-storage resource from which a new volume is allocated. A file-storage service class defines requirements for the file storage resources from which a new share is allocated.

You can associate a service class with one or more capacity pools. If you do, future provisioning requests can be satisfied only by storage resources in the specified capacity pools.

You can selectively grant users who do not have administrator privileges permission to provision storage by using the service class.

### Restrictions:

- To create a block-storage service class, you must have the IBM Tivoli Storage Productivity Center Advanced license.
- To associate the service class with a capacity pool, one or more capacity pools must be defined in the Tivoli Storage Productivity Center database. If no capacity pool is defined, the Specify Capacity Pool page of the wizard is not displayed.

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics > Cloud Configuration > Work With Service Classes**. To create a service class, complete the following steps:

1. From the Service Classes page, click **Create Service Class**. The Create Service Class wizard opens.
2. Select the type of service class that you want to create:
  - To create a block-storage service class, click the **Block** icon.
  - To create a file-storage service class, click the **File** icon.
3. In the Define Properties page, complete the following steps:
  - a. Specify a name for the service class.
  - b. Optional: Specify a service class description. Descriptions are displayed in the table of service classes on the Service Classes page.
  - c. Specify the general attributes of the service class. Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

**Tip:** To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.

4. Optional: Specify or modify the advanced properties of the service class.
  - a. From the Define Properties page, click **Advanced**.

- b. If you are creating a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
  - c. If you are creating a block-storage service class, select how the multipath driver on the host is configured. For hypervisors, the multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for Tivoli Storage Productivity Center.
  - d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
5. Optional: In the Specify Capacity Pool page, associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools.
  6. Optional: Grant permission to individual users, who do not have administrator privileges, to provision storage by using the service class. By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.
  7. Click **Finish**. The new service class is listed in the table of service classes on the Service Classes page, and can be specified in future provisioning requests.

#### Viewing the details of service classes:

You can view the provisioning requirements and storage constraints that a service class defines. You can also view the volumes or shares that were provisioned by using the service class.

From the Service Classes page, you can open a properties page to view the following information for a service class:

- The properties of a service class, such as the required RAID level or storage tier for a block-storage service class.
- The storage constraints that are imposed by the service class. The service class can allow provisioning from all available storage, or can limit provisioning to one or more capacity pools.
- The users who can provision storage by using the service class.
- The volumes or shares that were provisioned by using the service class.
- The candidate storage for the service class. The *candidate storage* is the set of storage resources that can provide the service class.

To open the service class properties window from the Service Classes page, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Service Classes**. The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions > View / Modify**. The service class properties window is displayed.

## Modifying service classes:

You can modify the attributes of a service class to change the provisioning requirements, capacity pool storage constraints, user permissions, and service class name.

You must have Administrator privileges to modify a service class.

**Important:** When volumes or shares are provisioned by specifying a service class, they are associated with the service class. The Volumes page and the Shares page in the web-based GUI can display the name of the service class that was used to provision the volume or share. If you modify a service class, the volumes or shares are still associated with the service class, but might be on storage resources that no longer satisfy the requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the new volumes or shares have attributes that they do not possess.

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics > Cloud Configuration > Work With Service Classes**. To modify a service class, complete the following steps:

1. From the list of service classes, select a service class and then select **Actions > View / Modify**. The service class properties window is displayed.
2. Optional: Click the **General** tab to modify the general attributes of the service class. Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

**Tip:** To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.

3. Optional: Specify or modify the advanced properties of the service class.
  - a. From the **General** tab, click **Advanced**.
  - b. If you are modifying a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
  - c. If you are modifying a block-storage service class, select how the multipath driver on the host is configured. For hypervisors, the multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for Tivoli Storage Productivity Center.
  - d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
4. Optional: Modify storage constraints. If one or more capacity pools are defined in the Tivoli Storage Productivity Center database, you can associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools. If a capacity pool was previously associated with the service class, you can disassociate it from the service class.
5. Optional: Modify user permissions. By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to

provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.

If users in the Monitor or External Application roles have permission to provision by using the service class, you can revoke that permission. You can specify that only administrators can provision by using the service class, or you can selectively remove users. To remove a user, right-click the user name in the table and then select **Remove** from the context-sensitive menu.

### Deleting service classes:

You can delete a service class that is no longer needed.

You must have Administrator privileges to delete a service class.

Deleting a service class does not affect the storage volumes or NAS shares that are already provisioned by using the service class. However, the storage volumes or NAS shares are no longer associated with any service class. When a volume that is not associated with a service class is optimized by the Optimization wizard, the volume can be migrated to any available storage pool. The storage pool to which the volume is migrated might not satisfy the requirements of the service class that was specified when the volume was provisioned.

You delete a service class from the Service Classes page.

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Service Classes**. The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions > Delete**.

### Tagging resources to satisfy custom requirements:

A service class can specify tags to create custom requirements for provisioning. Add the same tags to storage resources that satisfy the custom requirements. When the service class is specified during provisioning, only the tagged resources are candidates for provisioning.

You must have Administrator privileges to tag storage resources.

When you provision storage, you indicate your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have to be a candidate for provisioning. In addition to the default properties, a service class can define custom requirements by specifying up to three custom tags. To provide the service class, a storage resource must have all the same tags as the service class. The ability to define custom requirements by tagging service classes and storage resources gives you more control over the selection of storage resources during provisioning. By using tags, you can filter the candidates for provisioning by using any criteria that you choose.

When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Tivoli Storage Productivity Center identifies a storage pool for the volume. If the service class specifies tags, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, Tivoli Storage Productivity Center identifies a file system or Network Shared Disk (NSD) for the share. If the service class specifies tags, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

Tag only the storage resources that satisfy the intended custom requirements of the service class. During provisioning, only the tagged resources can provide the service class.

1. Open the resource list page for the type of storage resource that you want to tag. To satisfy the custom tag requirements of a block-storage service class, you can tag a storage pool or its containing block storage system. To satisfy the custom tag requirements of a file-storage service class, you can tag a file system or NSD, or the containing file storage system.

**Important:** The tags on a containing storage system apply to an internal pool, file system, or NSD resource only if the internal resource is not tagged.

- If you are tagging a storage pool, from the menu bar, go to **Storage > Pools**.
- If you are tagging the block storage system that contains the pools, complete the following steps:
  - a. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
  - b. Click the **Block Storage** tab.
- If you are tagging a file system or NSD, complete the following steps:
  - a. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
  - b. Click the **File Storage** tab.
  - c. From the list of storage systems, select the storage system that contains the file system or NSD, and then select **View Details**. Alternatively, you can select **Actions > View Details**.
  - d. In the Internal Resources area, click **File Systems** or **Network Shared Disks**.
- If you are tagging the file storage system that contains the file systems or NSDs, complete the following steps:
  - a. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
  - b. Click the **File Storage** tab.

2. In the table of storage systems, pools, file systems, or NSDs, select the storage resource that you want to tag, and then select **View Properties**. Alternatively, you can select **Actions > View Properties**.
3. In the properties notebook for the storage system, pool, file system, or NSD, click **Edit**.
4. In the custom tag fields of the properties notebook, specify tags that match the tags that are specified in the service class. To provide the service class, the resource must have all of the same tags.
5. In the properties notebook, click **Save**.

#### Setting access path host replacements:

If you are going to use the VMware vSphere Web Client extension to provision file shares, verify whether you must provide Tivoli Storage Productivity Center with an access path host replacement.



You need to provide an access path host replacement if the cluster name of an IBM SONAS or Storwize V7000 Unified is not a valid host name that is registered with a Domain Name System (DNS) server.

Complete the following steps to verify cluster host names and to provide access host name replacements where necessary:

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Click the **File Storage** tab.
3. For each IBM SONAS and Storwize V7000 Unified system, complete the following steps to determine whether you must specify a replacement host name.
  - a. From the **File System** tab of the Storage System page, select the storage system, and then select **Actions > View Details**.
  - b. From the list of File Internal Resources, click **Clusters**.
  - c. Determine whether the cluster name is a valid host name that is registered with a DNS server. If the cluster name is registered with a DNS server, the access path that is returned by the Tivoli Storage Productivity Center will be valid. If the cluster name is not valid, the VMware vSphere Web Client extension will not be able to access the share unless you specify a host name replacement for the access path.
4. For each IBM SONAS and Storwize V7000 Unified system that requires a host name replacement, complete the following steps:
  - a. From the **File Storage** tab of the Storage Systems page, select the storage system, and then select **Actions > View Properties**.
  - b. In the properties notebook for the storage system, click **Edit**.
  - c. In either the Custom Tag 1, Custom Tag 2, or Custom Tag 3 field, enter a fully qualified host name that is mapped to one or more interface nodes of the storage system.

**Requirement:** For all storage systems that require a replacement host name, ensure that you set the replacement host name by using the same custom tag. For all IBM SONAS and Storwize V7000 Unified systems that do not require a replacement host name, ensure that the custom tag is empty.

5. When you are creating or modifying file-storage service classes that might be used by the VMware vSphere Web Client extension, set the access host replacement to the custom tag that is set on the file storage systems. To set the access host replacement, complete these steps:
  - a. From the list of general properties in the Create Service Class wizard or the service class properties window, click **Advanced**.
  - b. On the **General** tab of the Advanced Properties window for the file-storage service class, set the access path host replacement to the custom tag that specifies the replacement host name. When the share is provisioned, the access path that is returned by Tivoli Storage Productivity Center includes the replacement host name if one is specified on the custom tag for the storage system. If the storage system does not specify a value for the custom tag, the cluster name is used in the access path.

## Configuring capacity pools

Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

## Adding resources to capacity pools:

Divide storage resources into capacity pools to later restrict provisioning or optimization to a specific set of resources. You can also use capacity pools to track the used and available space for block and file storage on any set of storage resources. You can add resources to an existing capacity pool, or you can create a new capacity pool.

You must have Administrator privileges to perform this task.

You can add storage resources to a capacity pool from resource list pages such as the Storage Systems page and the Pools page.

You can add the following resource types to a capacity pool:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM SONAS Network Shared Disks (NSDs)

When you add a storage system to a capacity pool, any storage pool, file system, or IBM SONAS NSD that is an internal resource of that storage system is also indirectly assigned to the capacity pool.

**Restriction:** A resource can be assigned to only one capacity pool. When you attempt to add one or more resources to a capacity pool that are already assigned to different capacity pool, IBM Tivoli Storage Productivity Center displays a message that lists the resources and their capacity pool assignments. When a message is displayed, ensure that you want to change the capacity pool assignments for the listed resources.

To add one or more resources to a capacity pool, complete these steps:

1. Open the appropriate list page for the resource you want to add. You can add an entire storage system, or you can add individual pools, file systems, or IBM SONAS NSDs.
  - To add a storage system, in the menu bar, go to **Storage > Storage Systems**.
  - To add a storage pool, in the menu bar, go to **Storage > Pools**.
  - To add a file system or IBM SONAS NSD, complete the following steps:
    - a. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
    - b. Click the **File Storage** tab.
    - c. Select the storage system that contains the file system or IBM SONAS NSD, and then select **View Details** from the context-sensitive menu.
    - d. From the Internal Resources list, click **File Systems** or **Network Shared Disks**.
2. To add one or more resource to a capacity pool, select the resources in the list, and then select **Actions > Add to Capacity Pool**. If one or more capacity pools are defined, a list of capacity pools is displayed. If no capacity pools are defined, the Create Capacity Pool window is displayed so that you can add resources to a new capacity pool.

**Tip:** The list of capacity pools and the Create Capacity Pool window shows capacity information. For the list of capacity pools, the Capacity column shows a bar that summarizes the used and free space in each capacity pool. In the Create Capacity Pool window, a capacity bar summarizes the used and free

space of the resources that are being added to the new capacity pool. Move the mouse pointer over the bar to display a text summary and details of the capacity usage.

3. Take one of the following actions:
  - If the Create Capacity Pool window is displayed, go to step 4.
  - If a list of capacity pools is displayed, add the selected resources to an existing capacity pool or to a new capacity pool.
    - To add the selected resource or resources to an existing capacity pool, select the appropriate row in the list of capacity pools, and click **Save**.
    - To add the selected resource to a new capacity pool, click **Create Capacity Pool**. The Create Capacity Pool window is displayed.
4. In the Create Capacity Pool window, specify a name for the capacity pool.
5. Optional: Specify a capacity pool description. Capacity pool descriptions appear in the list of capacity pools on the Capacity Pools page.
6. Optional: Specify custom tags to associate any significant information with the capacity pool. You can specify up to three custom tags. The custom tags can be displayed in the list of capacity pools on the Capacity Pools page.
7. Save the new capacity pool by clicking **OK**.

#### Viewing the details of capacity pools:

You can view details of a capacity pool, such as its total capacity, used and available space, and member resources.

The Capacity Pools page lists all of the capacity pools that are defined in IBM Tivoli Storage Productivity Center. For any capacity pool in the list, you can open a window to view more details about the capacity pool and its storage resources.

To view details of a capacity pool, complete these steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions > View / Modify**. A window is displayed, showing the following details:
  - A bar summarizing the used and free space in the capacity pool in gibibytes. The bar represents the total capacity of the capacity pool. The bar is divided into sections that represent the following measurements:
    - The amount of space that is used for file storage
    - The amount of available space for file storage
    - The amount of space that is used for block storage
    - The amount of block storage space that is not used for volumes, and is not reserved by pending or scheduled provisioning tasks.
    - The amount of block storage space that is reserved by pending or scheduled provisioning tasks

Move the mouse pointer over the bar to display a text summary and details of the capacity usage, and a legend for interpreting the colors on the bar.

- A table that lists the storage systems, storage pools, file systems, and IBM SONAS Network Shared Disks (NSDs) that are in the capacity pool. To view details of any of these resources, select the resource, and then select **Actions > View Properties**.

## Modifying capacity pools:

You can modify the attributes of a capacity pool such as its name and description, and you can remove storage resources from a capacity pool.

You must have Administrator privileges to perform this task.

To modify a capacity pool, complete these steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions > View / Modify**. A window that shows details of the capacity pool is displayed. A table lists the storage systems, storage pools, file systems, and IBM SONAS Network Shared Disks (NSDs) that are in the capacity pool.
3. Optional: Modify attributes of the capacity pool.
4. Optional: Remove one or more resources from the capacity pool. Removing a storage resource from a capacity pool does not delete the storage resource, but does remove its capacity from the capacity pool. To remove resources, select the appropriate resources in the list, and then select **Actions > Remove**. The resources are removed from the list, but are not removed from the capacity pool until you save your changes.
5. Click **Save**.

## Deleting capacity pools:

You can delete a capacity pool that is no longer needed.

You must have Administrator privileges to perform this task.

When you delete a capacity pool, the storage resources that were added to the capacity pool are no longer members of any capacity pool. Deleting a capacity pool does not affect any volumes or NAS shares that were provisioned from the capacity pool. However, the volumes or NAS shares are no longer associated with a capacity pool.

To delete a capacity pool, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions > Delete**
3. In the confirmation window, click **Delete**.

## Block storage: Calculating available space and determining the placement of volumes

For each block-storage service class, IBM Tivoli Storage Productivity Center calculates the amount of available space for volumes of that service class. When you request volumes of a particular service class, Tivoli Storage Productivity Center recommends the best location for the volumes from the available space.

The Service Classes page in the web-based GUI displays the available space for each block-storage service class. The available space is the amount of space that is available across storage pools for new volumes of the service class. Typically, not

all pools are used in the available-space calculation. Some pools might be eliminated because of the consolidated status of the pool or its storage system. Other pools might be eliminated because they do not satisfy the requirements of the service class. The set of storage pools that are used to calculate the available space are also the candidates for provisioning by using the service class.

When you request storage volumes by using the Provision Storage wizard, you specify the capacity that is required for one or more volumes, and you specify a service class. You can also specify a capacity pool from which the volumes must be provisioned. Based on the volume size that is requested and if a capacity pool is specified, the candidates for provisioning are further refined to include only the pools that have enough available space and are in the specified capacity pool. From the remaining candidates for provisioning, Tivoli Storage Productivity Center creates a provisioning task that identifies the best location for the new volumes.

### **How available space for a block-storage service class is calculated**

For each block-storage service class, Tivoli Storage Productivity Center calculates the available space for volumes of that service class. Only storage pools that are candidates for provisioning by using the service class are used to calculate the available space.

Identifying candidates for provisioning is a filtering process. For example, if the service class specifies that RAID 10 is required, then any storage pool that is not configured at RAID 10 is eliminated as a candidate for provisioning. Tivoli Storage Productivity Center filters the set of available storage resources against certain properties of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning. Only certain properties of a service class are used to identify candidates for provisioning. More advanced properties contain configuration instructions, and are not used to identify candidates for provisioning.

The candidates for provisioning are identified by eliminating storage pools that cannot provide the service class, as follows:

- If a storage pool is a primordial pool or if it is formatted for count key data (CKD) volumes, it is not a candidate for provisioning.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal storage pools are candidates for provisioning. If a storage pool has a consolidated status of Error, it is not a candidate for provisioning.

**Remember:** A consolidated status is different from a propagated status, which is shown in the web-based GUI. A consolidated status is the status reported by a storage resource without considering the status of its internal resources. A propagated status is the combined status of the monitored, internal resources for the higher-level resource, and the consolidated status of the higher-level resource itself. The status for a storage system that is shown on the Storage Systems page and the status for a pool that is shown on the Pools page is the propagated status. The propagated status might not have the same value as the consolidated status, which is used to eliminate pools as candidates for provisioning.

Although the propagated status of a resource might not have the same value as its consolidated status, the propagated status is always at least as severe as the consolidated status. The propagated status reflects the most severe consolidated

status of any internal resource or the consolidated status of the resource itself. If the propagated status of a resource is Normal, then its consolidated status is also Normal.

**Tip:** The Service Classes page in the web-based GUI displays both the available space and the unavailable space for a service class. The Unavailable Space column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

- If a service class allows provisioning only from certain capacity pools, any storage pool that is not in a specified capacity pool is eliminated as a candidate for provisioning.

**Tip:** You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.

- A storage pool must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect whether the storage pool is a candidate for provisioning.

*Table 71. Filtering candidates for provisioning based on service class properties*

Service class property:	Candidates for provisioning:
Storage tier	If specified in the service class, storage pools that are assigned to any storage tier within the specified range are candidates for provisioning.
RAID level	If a RAID level is specified, only storage pools that are configured at the specified RAID level are candidates for provisioning. <b>Exception:</b> The IBM XIV Storage System uses a nontraditional approach to data redundancy and does not support any of conventional RAID levels. Tivoli Storage Productivity Center always considers XIV system storage pools to be configured at RAID 10. <b>Restriction:</b> Storage pools that support multiple RAID levels are not candidates for provisioning when a service class specifies a RAID level. Storage pools that support multiple RAID levels are not candidates even when the specified RAID level is one of the RAID levels that are supported by the storage pool. Storage pools that support multiple RAID levels are only candidates for provisioning when the service class specifies that any RAID level is allowed.
Virtualization	If virtualization is required by the service class, only storage pools on IBM storage virtualizers are candidates for provisioning. The IBM storage virtualizers are SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified  If virtualization is not allowed by the service class, only storage pools that are not on IBM storage virtualizers are candidates for provisioning.
Encryption	If encryption is required by the service class, only storage pools on DS8000 storage systems that are encrypted and configured to use the encryption group that is specified in the service class are candidates for provisioning.  If encryption is not allowed by the service class, only pools that are not encrypted are candidates for provisioning.

Table 71. Filtering candidates for provisioning based on service class properties (continued)

Service class property:	Candidates for provisioning:
Thin provisioning	<p>If thin provisioning is required by the service class, only storage pools that are configured for thin provisioning are candidates for provisioning. Thin provisioning is supported by the XIV system, DS8000, SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified. During provisioning, a thin-provisioned volume is created.</p> <p>If thin provisioning is not allowed by the service class, any storage pool that is not thin provisioned is a candidate for provisioning. In addition, any DS8000, SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified storage pool is a candidate for provisioning, regardless of whether it is configured for thin provisioning. During provisioning, the volume that is created is not thin provisioned.</p>
Compression	<p>If compression is required by the service class, at least one I/O group of an IBM storage virtualizer must be enabled for compression. Only pools of the IBM storage virtualizer are candidates for provisioning.</p>
Overallocation limit	<p>If thin provisioning is required by the service class, Tivoli Storage Productivity Center identifies the storage pools that can thin provision volumes as candidates for provisioning. If an overallocation limit is specified, Tivoli Storage Productivity Center refines that list to include only the storage pools that have an overallocation percentage that does not exceed the overallocation limit.</p> <p>The overallocation percentage for a pool is determined by using the following measurements, as shown on the Pools page:</p> <ul style="list-style-type: none"> <li>• For an XIV system storage pool, the Soft Space divided by the Capacity.</li> <li>• For other storage system pools, the Virtual Allocation.</li> </ul>
Resource Tags	<p>Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.</p>

After the storage pools that can provide the service class are identified, Tivoli Storage Productivity Center calculates the amount of space that is available for provisioning new volumes of the service class. The calculated volume capacity is shown in the Available Space column on the Service Classes page.

The available-space calculation differs depending on whether the block-storage service class requires thin provisioning, as follows:

- If thin provisioning is not allowed, the available space is the sum of the following measurements, as shown on the Pools page:
  - For each XIV System pool, the Available Soft Space minus any Reserved Pool Space
  - For other storage system pools, the Available Pool Space minus any Reserved Pool Space.
- If thin provisioning is required, the available space is the total of the following measurements, as shown on the Pools page, and calculations:
  - For each XIV System pool, the Available Soft Space minus any Reserved Pool Space

- For other storage system pools, the pool Capacity multiplied by the overallocation limit that is specified in the service class, minus the Total Volume Capacity. Any Reserved Pool Space is subtracted from the result. For DS8000 storage pools, any space that is used by Track Space Efficient (TSE) volumes is subtracted from the calculated available space.

## How the recommended placement of storage volumes is determined

When you request one or more new volumes by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new volumes. The same filtering that was used to determine the available space for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of storage pools from which the storage can be provisioned. This set can be all the storage pools that are known to Tivoli Storage Productivity Center, or can be limited to the storage pools in a particular capacity pool. Only the storage pools in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available space for the service class.
- Any storage pool that cannot provide the requested volume capacity is eliminated as a candidate for provisioning.
- If virtualization and VDisk mirroring are specified by the service class, then a VDisk copy of the volume must be made in a second storage pool on the same IBM storage virtualizer. If the storage pools on the IBM storage virtualizer cannot provide the requested capacity for both the volume and its mirrored copy, the pools are eliminated as candidates for provisioning.
- If an overallocation limit is specified in the service class, only the storage pools that will not exceed the limit after the new volume or volumes are provisioned are candidates for provisioning.

**Remember:** Storage pools that have an overallocation percentage that does not exceed the overallocation limit, are included in the available-space calculation. However, if the overallocation percentage would exceed the overallocation limit after the requested volume space is included, the pools are not candidates for provisioning.

After the storage pools that can satisfy the provisioning request are identified, Tivoli Storage Productivity Center finds the best location for the storage. The best location for the storage is based on the unallocated volume space in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is given to systems that have available performance data.

## File storage: Calculating available space and determining the placement of shares

For each file-storage service class, IBM Tivoli Storage Productivity Center calculates the amount of available space for NAS file shares of that service class. When you request file shares of a particular service class, IBM Tivoli Storage Productivity Center recommends the best location for the shares from the available space.

For each file-storage service class, Tivoli Storage Productivity Center calculates the amount of available space for NAS file shares of that service class. When you request file shares of a particular service class, Tivoli Storage Productivity Center recommends the best location for the shares from the available space.



The Service Classes page in the web-based GUI displays the available space for each file-storage service class. The available space is the amount of space that is available across all IBM SONAS and IBM Storwize V7000 Unified file systems and unused IBM SONAS system NSDs that are known to Tivoli Storage Productivity Center, and that can satisfy the requirements of the service class. The set of file storage resources that are used to calculate the available space are also the candidates for provisioning by using the service class.

When you request a file share by using the Provision Storage wizard, you specify the capacity that is required for the share, and you specify a service class. You can also specify a capacity pool from which shares must be provisioned. Based on the requested size of the share and whether a capacity pool is specified, the candidates for provisioning are further refined to include only the file storage resources that have enough available space and are in the specified capacity pool. From the remaining candidates for provisioning, Tivoli Storage Productivity Center creates a provisioning task that identifies the best location for the new share.

### **How available space for a file-storage service class is calculated**

For each file-storage service class, Tivoli Storage Productivity Center calculates the available space for file shares of the service class. Only the file storage resources that are candidates for provisioning by using the service class are used to calculate the available space.

Identifying candidates for provisioning is a filtering process. For example, if the service class requires dedicated storage, then the share must be created on an unused NSD within an IBM SONAS system. File systems that could be candidates for provisioning if shared storage is allowed are eliminated as candidates for provisioning when dedicated storage is required. Tivoli Storage Productivity Center filters the set of available storage resources against each attribute of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning.

The candidates for provisioning are identified by eliminating storage resources that cannot provide the service class, as follows:

- Only IBM SONAS and IBM Storwize V7000 Unified systems are supported for provisioning file shares. No other storage system type is a candidate for provisioning file shares.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal NSDs or file systems are included in the available-space calculation.
- If an NSD or file system has a consolidated status of Error, it is not included in the available-space calculation.

**Tip:** The status of file systems is not displayed in the web-based GUI. View the status on the NSD on which a file system resides to determine the file system status.

- If the storage volume or NSD on which a file system resides has a consolidated status of Error, the file system is not included in the available-space calculation.

**Remember:** A consolidated status is different from a propagated status, which is shown in the web-based GUI. A consolidated status is the status reported by a storage resource without considering the status of its internal resources. A propagated status is the combined status of the monitored, internal resources for the top-level resource. The status for a storage system that is shown on the Storage Systems page, the status for volume that is shown on the Volumes page,

and the status of an NSD that is shown on the storage system details page is the propagated status. The propagated status might not have the same value as the consolidated status, which is used to eliminate file storage resources as candidates for provisioning.

Although the propagated status of a resource might not have the same value as its consolidated status, the propagated status is always at least as severe as the consolidated status. The propagated status reflects the most severe consolidated status of any internal resource or the consolidated status of the resource itself. If the propagated status of a resource is Normal, then its consolidated status is also Normal.

**Tip:** The Service Classes page in the web-based GUI displays both the available space and the unavailable space for a service class. The Unavailable Space column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

- If a service class allows provisioning only from certain capacity pools, any storage resource that is not in a specified capacity pool is eliminated as a candidate for provisioning.

**Tip:** You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.

- A file system or NSD must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect the storage placement determination.

*Table 72. Filtering candidates for provisioning based on service class properties*

Service class attribute:	Candidates for provisioning:
Shared storage / Dedicated storage	If the service class specifies dedicated storage, the share can be created only on an unused NSD within an IBM SONAS system. File systems and other file storage resources are eliminated as candidates for provisioning.
Resource tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

After the file storage resources that can provide the service class are identified, Tivoli Storage Productivity Center calculates the amount of space that is available for providing new files shares of the service class.

The available-space calculation differs depending on whether the file-storage service class requires dedicated storage or allows shared storage, as follows:

- If the file-storage service class requires dedicated storage, the available space is the total available space on unused NSDs within each IBM SONAS system. You can verify that an NSD is unused, and view its available space, from the details page of an IBM SONAS system. To view information about the NSDs that are associated with an IBM SONAS system, click **Network Shared Disks**. If the

NSD is unused, the File Systems column shows 0. The Available Space column shows the amount of space on the NSD, which is the measurement that is used in the available-space calculation for the service class.

- If the file-storage service class allows shared storage, the file system in which the NAS share is provisioned can contain other NAS shares. In this case the available space is the total of the following measurements:
  - The available space on unused NSDs within each IBM Scale Out Network Attached Storage (IBM SONAS) system.
  - The available space for each file system within each IBM SONAS and IBM Storwize V7000 Unified storage system. You can view the file system available space from the details page of a file storage system. To view information about the file systems that are in a storage system, click **File Systems**. The Available Space column shows the amount of available space on each file system, which is the measurement that is used in the available-space calculation for the service class.

### How the recommended placement of file shares is determined

When you request a file share by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new share. The same filtering that was used to determine the available space for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of file storage resources from which the storage can be provisioned. This set can be all the storage resources that are known to Tivoli Storage Productivity Center, or can be limited to the storage resources in a particular capacity pool. Only the storage resources in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available space for the service class.
- Any storage resource that cannot provide the requested capacity for the share is eliminated as a candidate for provisioning.

The shares are created on file systems or IBM SONAS NSDs. After the file systems and NSDs that can satisfy the provisioning request are identified, Tivoli Storage Productivity Center identifies the best location for the storage. The best location for the storage is based on the available space on the file system or NSD.

## Setting the zoning policy

You can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, Tivoli Storage Productivity Center can create zones during storage provisioning to connect a server to a storage system.

When Tivoli Storage Productivity Center creates a provisioning task, it identifies the best location for the new storage that satisfies the requirements of the service class. If automatic zoning is enabled, then, during provisioning, existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more zones are created between a host initiator port and a controller, node, or module port.

To set or modify the zoning policy, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics > Provisioning**, and click **Set Zoning Policy**.
2. Enable or disable automatic zoning:

- Enable automatic zoning to allow Tivoli Storage Productivity Center to create zones during storage provisioning. The naming convention for new zones is *host\_storage-system\_suffix-number*. When you enable automatic zoning, you can specify the following options:

**Zone name prefix**

If you specify a zone name prefix, all zones that are automatically created by Tivoli Storage Productivity Center are prefixed with a string you specify. This prefix can help you identify which zones were automatically created, and which were created manually.

**Make changes to the active zone set**

Whether changes are made to the active zone set, or to a new inactive active zone set. If this check box is selected, changes are made to the active zone set. If the check box is cleared, changes are made to the inactive zone set. The new inactive zone will contain only the new zones.

- Disable automatic zoning if you want Tivoli Storage Productivity Center to use only existing zones.

3. Click **Save**.

## Provisioning storage with the Tivoli Storage Productivity Center GUI

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

The prerequisites for provisioning storage are as follows:

- You must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.
- You must have Administrator privileges or permission to use one or more of the service classes that are created for block-storage or file-service provisioning.
- You must add and probe block-storage or file-storage resources, or both.

The following restrictions apply to volume provisioning:

- If the service class that is associated with the provisioning request requires multipathing or if automatic zoning is enabled, the fabrics and switches must be managed by Tivoli Storage Productivity Center. If the fabric is not managed by Tivoli Storage Productivity Center, the fabric-related configuration options are ignored when the provisioning task is executed.
- If you provision volumes to physical servers and servers that are virtual machines, you must create separate provisioning tasks with the Provision Storage wizard.
- If you provision multiple servers with storage volumes, the servers must all run on the same operating system. If the servers are virtual machines, the hypervisors that manage the servers must all run on the same operating system.
- If the provisioning request involves assigning multiple volumes to a resource, the volumes are sourced from pools in the same storage system.

NAS file shares can be provisioned only from IBM SONAS or Storwize V7000 Unified storage systems.

In the Provision Storage wizard, specify the capacity that is required and the service class that you want to associate with the provisioning task. For shares, you

must also specify at least one file access protocol. Tivoli Storage Productivity Center identifies the storage resources that can provide the storage capacity and that meet the criteria that are specified for the service class. From the set of storage resources that can provide the capacity and service class, Tivoli Storage Productivity Center identifies the best location for the storage and creates a provisioning task. Depending on the privileges that are associated with your role, you can save, execute, or schedule the provisioning task.

## Provisioning volumes to servers

To provision volumes to one or more servers, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity and specifying requirements by using a service class.

You can provision volumes to one or more servers. If a server is a virtual machine, the storage is provisioned to the hypervisor that is managing the server.

1. In the menu bar in the web-based GUI, go to **Servers > Servers..**

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Servers**.

2. Select one or more servers, and then select **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the storage volumes.

The Provision Storage wizard creates a provisioning task.

### Related tasks:

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

## Provisioning volumes to hypervisors

To provision volumes to one or more hypervisors, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity and specifying requirements by using a service class.

**Restriction:** The multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for Tivoli Storage Productivity Center.

You can provision volumes to one or more hypervisors.

1. **Servers > Hypervisors**

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Hypervisors**.

2. Select one or more hypervisors, and then click **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the storage volumes.

The Provision Storage wizard creates a provisioning task.

### Related tasks:

“Provisioning block storage” on page 267

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

## Provisioning volumes to a server cluster

To assign volumes to a cluster of servers, use the Provision Storage wizard.

If the cluster is on servers that are virtual machines, the storage is assigned to the hypervisors that manage the servers. The volumes are assigned to all of the nodes in the cluster.

1. In the menu bar in the web-based GUI, go to **Servers > Servers**, and click **Clusters**.

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Servers**, and click **Clusters**.

2. Select a cluster, and then click **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the storage volumes.

The Provision Storage wizard creates a provisioning task.

### Related tasks:

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

## Provisioning volumes to a hypervisor cluster

To assign volumes to a cluster of hypervisors, use the Provision Storage wizard.

**Restriction:** The multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for Tivoli Storage Productivity Center.

The volumes are assigned to all of the nodes in the hypervisor cluster.

1. In the menu bar in the web-based GUI, go to **Servers > Hypervisors**, and click **Clusters**.

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Hypervisors** and click **Clusters**.

2. Select a cluster, and then click **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the storage volumes.

The Provision Storage wizard creates a provisioning task.

### Related tasks:

“Provisioning block storage” on page 267

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

### Provisioning shares to servers

To provision a network-attached storage (NAS) file share to one or more servers, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You can provision a NAS file share to one or more servers. The share is made available to the selected servers, but is not automatically mounted from the selected servers.

1. In the menu bar in the web-based GUI, go to **Servers > Servers..**

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Servers**.

2. Select one or more servers, and then select **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the share.

The Provision Storage wizard creates a provisioning task.

#### Related tasks:

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

### Provisioning shares to hypervisors

To provision a network-attached storage (NAS) file share to one or more hypervisors, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You can provision a NAS file share to one or more hypervisors. The share is made available to the selected hypervisors, but is not automatically mounted from the selected hypervisors.

1. **Servers > Hypervisors**

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, and click **Provision to Hypervisors**.

2. Select one or more hypervisors, and then select **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the share.

The Provision Storage wizard creates a provisioning task.

#### Related tasks:

“Provisioning file storage” on page 268

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

## Provisioning shares to a server cluster

To assign a network-attached storage (NAS) file share to a server cluster, use the Provision Storage wizard.

You can provision a NAS file share to a server cluster. The share is made available to the servers in the cluster, but is not automatically mounted from the selected servers.

1. In the menu bar in the web-based GUI, go to **Servers > Servers**, and click **Clusters**.

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Servers**, and click **Clusters**.

2. Select a cluster, and then select **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the share.

The Provision Storage wizard creates a provisioning task.

### Related tasks:

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

## Provisioning shares to a hypervisor cluster

To assign a network-attached storage (NAS) file share to a hypervisor cluster, use the Provision Storage wizard.

You can provision a NAS file share to a hypervisor cluster. The share is made available to the hypervisors in the cluster, but is not automatically mounted from the selected hypervisors.

1. In the menu bar in the web-based GUI, go to **Servers > Hypervisors**, and click **Clusters**.

**Tip:** Alternatively, you can go to **Advanced Analytics > Provisioning**, click **Provision to Hypervisors** and click **Clusters**.

2. Select a cluster, and then select **Actions > Provision Storage**.
3. Follow the instructions in the Provision Storage wizard to assign the share.

The Provision Storage wizard creates a provisioning task.

### Related tasks:

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

## Provisioning storage with the vSphere Web Client extension

You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage and monitor your virtual environment and storage devices.

You can provision block and file storage to hypervisors from storage devices that are monitored by Tivoli Storage Productivity Center. You can monitor the progress of your storage provisioning activities in task information in the vSphere Web Client.



## Provisioning block storage

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

Ensure that you complete the following tasks in Tivoli Storage Productivity Center:

- If you want to provision storage by using service classes to describe performance and storage requirements and to implement multipathing policies, ensure that the service classes are created in Tivoli Storage Productivity Center. Tivoli Storage Productivity Center provides a number of predefined service classes, which you can use, modify, or delete.
- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in Tivoli Storage Productivity Center.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into Tivoli Storage Productivity Center. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by Tivoli Storage Productivity Center.
- Ensure that the hypervisor has fabric connectivity. If automatic zoning is enabled, the fabrics and switches must be managed by Tivoli Storage Productivity Center.

For more information about configuring for storage cloud provisioning, see the Tivoli Storage Productivity Center information center. Search for *Configuring for storage cloud provisioning*.

Start the vSphere Web Client.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select the hypervisor for which you want to provision storage.
3. Click the **Actions** menu item, and select **All TPC Actions**.
4. To open the **Provision Block Storage** wizard, select **Provision Block Storage...**
5. In the wizard, enter a size for the total usable storage space on the volume, and select the unit of measurement from the list.
6. Select the service class for the block storage that you want to associate with the volume. You can select only service classes with available capacity that you are authorized to use.
7. Select the capacity pool from which you want the volume to be provisioned. You can select only capacity pools with available capacity.
8. Optional: If you want to create a data store on the storage volume when it is provisioned, select the **Create datastore** check box, and enter a name for the data store. The minimum volume size for creating a data store is 2 GiB.
9. Click **OK** to provision the volume.
10. To monitor the progress of the provisioning task, click the **Monitor** tab, and click the **Tasks** tab.
11. You can monitor the progress of the creation of the data store by clicking the **Monitor** tab and then clicking the **Tasks** tab. The data store creation task is called **Create VMFS datastore**.
12. You can monitor the progress of the multipathing task by clicking the **Monitor** tab and clicking the **Tasks** tab. The multipathing task is called **Set logical unit policy**.

**Related concepts:**

“Service classes” on page 70

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

“Capacity pools” on page 74

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

**Related tasks:**

“Provisioning file storage”

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

**Provisioning file storage**

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

Ensure that you complete the following tasks in Tivoli Storage Productivity Center:

- If you want to provision storage by using service classes to describe performance and storage requirements, ensure that the service classes are created in Tivoli Storage Productivity Center. Tivoli Storage Productivity Center provides a number of predefined service classes, which you can use, modify, or delete.
- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in Tivoli Storage Productivity Center.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into Tivoli Storage Productivity Center. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by Tivoli Storage Productivity Center.
- Determine whether you must set access path host name replacements. For more information, see “Special consideration for provisioning by using the VMware vSphere Web Client extension” on page 74.

For more information, see “Cloud configuration” on page 70.

Start the vSphere Web Client.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select the hypervisor for which you want to provision storage.
3. Click the **Actions** menu item, and select **All TPC Actions**.
4. To open the **Provision File Storage** wizard, select **Provision File Storage...**
5. In the **Share Requirements** page of the wizard, enter a name for the share.
6. Enter a size for the total usable storage space on the share, and select the unit of measurement from the list.
7. Select the service class for the file storage that you want to associate with the share. You can select only service classes with available capacity that you are authorized to use.
8. Select the capacity pool from which you want the share to be provisioned. You can select only capacity pools with available capacity.

9. Enter the name of the data store that you want to associate with the file system.
10. Click **Next** to continue.
11. In the **Share Permissions** page of the wizard, specify the user and group with permission to mount the share on the hypervisor.
12. Select the **Read-only** check box if you want to limit the user or group to viewing files on the share.
13. Select the **Secure** check box to create the share with the secure flag set.
14. Click **OK** to provision the share.
15. To monitor the progress of the provisioning task, click the **Monitor** tab, and click the **Tasks** tab.
16. You can also monitor the progress of the creation of the data store by clicking the **Monitor** tab and then clicking the **Tasks** tab. The data store creation task is called **Create NAS datastore**.

**Related concepts:**

“Service classes” on page 70

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

“Capacity pools” on page 74

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

**Related tasks:**

“Setting access path host replacements” on page 250

If you are going to use the VMware vSphere Web Client extension to provision file shares, verify whether you must provide Tivoli Storage Productivity Center with an access path host replacement.

“Provisioning block storage” on page 267

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

---

## Optimizing storage tiering

To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

Before you open the Analyze Tiering wizard, complete the following tasks:

- Probe the storage virtualizers
- Collect performance data
- Set the tier level of the storage pools that you analyze and that you select as target pools

The purpose of analyzing tiering is to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes based on the volume

workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated.

Storage costs can be reduced, for example, by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers. Storage performance can be improved and used more efficiently by moving volumes with heavy workloads to the tiers that meet the workload requirements of the volumes. When you implement the recommendations, you ensure that the volumes are placed on the tiers that best match the workload requirements of the volumes.

You can analyze the tiering of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

Volumes can be moved to tiered storage pools on the same storage virtualizer, but volumes cannot be moved from one storage virtualizer pool to another storage virtualizer pool.

**Optimizing volumes with service classes:** Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

**Related tasks:**

“Re-tiering volumes to balance pools” on page 95

In this scenario, you must re-tier some of the volumes to help resolve the performance hot spots that you identified. To balance and re-tier volumes, you run the Analyze Tiering wizard.

“Tiering volumes by I/O density and file usage” on page 97

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To ensure that only volumes that require the performance capability of tier 1 storage, you add minimum thresholds for I/O density and file usage for each storage tier.

“Tiering volumes by I/O rate” on page 99

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To restrict the placement of volumes in tier 1 storage pools, you add minimum threshold values for the I/O rate.

## Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Before you open the Analyze Tiering wizard or the Balance Pools wizard, you must set the tier level of pools.

To analyze tiering, you must set the tier level of the source pools and the target pools. The source pools are the pools that you select for analysis. The target pools are the pools that are used to relocate volumes to lower or higher tiers of storage. To balance pools, you must set the tier level of the pools that you want to balance.

If you do not set the tier level of the pools, you are prompted to enter the tier level on the More Information Required page of the Analyze Tiering or the Balance Pools wizard.

**Tip:** You can also set or modify the tier level of pools on the **General** tab of the View Properties page for the storage pool.

You can also assign tier levels to storage systems that are not storage virtualizers. However, you cannot balance the pools or analyze the tiering of pools other than storage virtualizer pools.

1. In the menu bar in the web-based GUI, go to **Storage > Pools**.
2. On the Pools page, select one or more storage pools.
3. Click **Set > Tier**, and then select a tier level. To remove the tier level, click **None**.

The tier level that you selected for the storage pools is shown in the **Tier** column.

## Analyzing tiering by servers

To analyze the tiering of volumes on storage virtualizers that are connected to servers, use the Analyze Tiering wizard.

The servers that you select must be connected to one or more of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. In the menu bar in the web-based GUI, go to **Servers > Servers**.
2. Click one or more servers, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions > Analyze Tiering**.

The Analyze Tiering wizard opens.

## Analyzing tiering by hypervisors

To analyze the tiering of volumes on storage virtualizers that are connected to hypervisors, use the Analyze Tiering wizard.

The hypervisors that you select must be connected to one or more of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. In the menu bar in the web-based GUI, go to **Servers > Hypervisors**.
2. Click one or more servers, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions > Analyze Tiering**.

The Analyze Tiering wizard opens.

## Analyzing tiering by storage virtualizers

To analyze the tiering of volumes on storage virtualizers, use the Analyze Tiering wizard.

You selected one or more of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
2. Click one or more storage virtualizers, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions > Analyze Tiering**.

The Analyze Tiering wizard opens.

## Analyzing tiering by storage pools

To analyze the tiering of volumes on storage pools, use the Analyze Tiering wizard.

The storage pools that you select are on a back-end storage system that is connected to one of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. In the menu bar in the web-based GUI, go to **Storage > Pools**.
2. Click one or more storage pools, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions > Analyze Tiering**.

The Analyze Tiering wizard opens.

## Analyzing tiering by volumes

To analyze the tiering of volumes, use the Analyze Tiering wizard.

The volumes that you select are in storage pools on back-end storage systems that are connected to one of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. In the menu bar in the web-based GUI, go to **Storage > Volumes**.
2. Click one or more volumes, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions > Analyze Tiering**.

The Analyze Tiering wizard opens.

## Balancing pools and re-tiering volumes

To re-tier volumes to balance pools across storage tiers, use the Analyze Tiering wizard.

1. Select the resources that you want to analyze.
2. Specify the target storage pools for the volumes.

**Mirrored volumes:** If you want to optimize mirrored volumes, select one or more target pools that do not contain either the primary volume or the secondary volume. Primary volumes and secondary volumes cannot be placed in the same destination pool.

3. Include or exclude volumes in mirrored volume relationships from the analysis.
4. On the Optimize the Placement of Volumes page, specify the number of days for estimating the activity of the pools.
5. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool. Recommendations are only generated to move one or more of the related volumes when all of the related volumes can be placed in the same destination storage pool.

**Multiple host connections to the same hypervisor or server:** If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

6. For each tier, specify the activity limit for the destination pools.
7. Click **Analyze**.

When you complete entering the information to run the Analyze Tiering wizard, recommendations for balancing the pools across each tier, and for re-tiering volumes on tiers that cannot be balanced are shown on the Tiering Analysis page. You can create a schedule to run the analysis that is based on the criteria that you selected in the Analyze Tiering wizard.

## Tiering volumes by I/O density, I/O rate, and file usage

To tier volumes by setting minimum thresholds for I/O rate or I/O density, and file usage, use the Analyze Tiering wizard.

1. Select the resources that you want to analyze. The source storage pools that are related to the resources that you selected are analyzed to determine whether they meet the workload requirements of the volumes. If the workload requirements of the volume in its current tier are not met, the volume is a candidate for relocation.
2. Specify the target storage pools for the volumes.

**Mirrored volumes:** If you want to optimize mirrored volumes, select one or more target pools that do not contain either the primary volume or the secondary volume. Primary volumes and secondary volumes cannot be placed in the same destination pool.

3. Include or exclude volumes in mirrored volume relationships from the analysis.

4. On the Optimize the Placement of Volumes page, specify the number of days for estimating the activity of the pools.
5. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool. Recommendations are only generated to move one or more of the related volumes when all of the related volumes can be placed in the same destination storage pool.

**Multiple host connections to the same hypervisor or server:** If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

6. Select **Show advanced options**.
7. To specify the activity limit and minimum threshold levels for placing volumes on each tier, complete these steps:
  - a. Set an activity limit. To ensure that the performance of the pools does not deteriorate when volumes are added, you set an activity limit for all of the pools on each tier.
  - b. Select **I/O Rate**, or **I/O Density**, and enter a minimum threshold level.
  - c. Select the percentage of files that must be accessed and the period for accessing the files.
8. Click **Analyze**.

When you complete entering the information to run the Analyze Tiering wizard, recommendations for tiering the volumes are shown on the Tiering Analysis page. You can create a schedule to run the analysis that is based on the criteria that you selected in the Analyze Tiering wizard.

## Modifying the criteria for analyzing tiering

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

You must run the Analyze Tiering wizard to create a tiering analysis task.

The criteria that you modify are used to analyze the source volumes and target volumes that were selected when you ran the Analyze Tiering wizard.

1. On the Navigation pane, click **Home > Tasks**.
2. Right-click a tiering analysis task, and click **Edit**.
3. Modify the tiering thresholds, and then click **Analyze**.

The source volumes are analyzed to determine whether one or more of the volumes require re-tiering based on the criteria that you specified. If one or more of the volumes require re-tiering, recommendations are generated to move the volumes to the target pools that meet the criteria that you specified.

## Storage tiering policies

In the Analyze Tiering wizard, you can tier volumes based on tiering criteria or enforce the balancing of pools by re-tiering volumes.



On the Optimize the Placement of Volumes page of the Analyze Tiering wizard, you can choose how you want to optimize the storage resources that you selected for analysis.

If you want to tier volumes based on the workload requirements of the volumes, you set minimum thresholds for each tier of storage. You can specify minimum thresholds for I/O density or I/O rate, and file usage. You can also combine minimum thresholds for I/O density and file usage, and for I/O rate and file usage. The workload of the volumes is evaluated to determine whether the volumes require re-tiering.

If you want to balance pools and restrict the re-tiering of volumes to pools that cannot be balanced across each tier, you do not specify tiering criteria. The process for balancing pools and re-tiering volumes is as follows:

1. Starting with the lowest tier of storage, the source pools are balanced.
2. If the pools cannot be balanced, the most-active volumes are moved to a higher tier of storage. This step does not apply to pools on the highest tier of storage.
3. The source pools are balanced again.
4. If the pools cannot be balanced, the least-active volumes are moved to a lower tier. This step does not apply to pools on the lowest tier of storage.
5. The source pools are balanced again.
6. Steps 1 - 5 are repeated until the highest tier is balanced.

Regardless which option that you choose to optimize storage, you must set an activity limit for each tier of storage. Activity limits are set to maintain the performance capability of the destination pools.

**Related tasks:**

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

“Re-tiering volumes to balance pools” on page 95

In this scenario, you must re-tier some of the volumes to help resolve the performance hot spots that you identified. To balance and re-tier volumes, you run the Analyze Tiering wizard.

“Tiering volumes by I/O density and file usage” on page 97

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To ensure that only volumes that require the performance capability of tier 1 storage, you add minimum thresholds for I/O density and file usage for each storage tier.

“Tiering volumes by I/O rate” on page 99

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To restrict the placement of volumes in tier 1 storage pools, you add minimum threshold values for the I/O rate.

**Related reference:**

“Criteria for identifying the pools that require balancing” on page 278

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

---

## Optimizing storage pools

Use the Balance Pools wizard to analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

Before you open the Balance Pools wizard, complete the following tasks:

- Probe the storage virtualizer. You must probe the storage virtualizer to collect information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.
- Collect performance data. To gauge the average activity level of the volumes and pools, you must collect sufficient performance data for the storage virtualizer.
- Assign a tier level to the pools that you want to balance. If you select pools that are not assigned a tier level, you are prompted to set the tier level on the More Information Required page of the Balance Pools wizard.

You can balance the workload of volumes in pools on the same tier on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

To balance pools, you must select two or more pools on the same tier level. For example, you select three tier 1 pools, and four tier 3 pools. The pools on tier 1 are analyzed and balanced, and the pools on tier 3 are analyzed and balanced.

**Optimizing volumes with service classes:** Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

### Related tasks:

“Setting the tier level of storage pools” on page 270

To analyze tiering and to balance pools, you must set the tier level of storage pools.

“Redistributing volume workloads to balance pools” on page 94

In this scenario, you identify performance hot spots and then run the Balance Pools wizard to help resolve them.

“Re-tiering volumes to balance pools” on page 95

In this scenario, you must re-tier some of the volumes to help resolve the performance hot spots that you identified. To balance and re-tier volumes, you run the Analyze Tiering wizard.

## Balancing pools

To distribute the workload of volumes across pools on the same tier, use the Balance Pools wizard. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier.

1. In the menu bar in the web-based GUI, go to **Storage > Pools**.
2. Click two or more storage pools with the same tier level, and then select **Balance Pools** from the context-sensitive menu. Alternatively, you can click **Actions > Balance Pools**.

**Tip:** If you want to analyze pools that contain volumes in mirrored volume relationships, select one or more pools that do not contain either the primary

volume or the secondary volume. The primary volume and the secondary volume cannot be placed in the same destination pool.

3. Choose whether to analyze volumes in mirrored volume relationships. These options are only available if you select one or more volumes with mirrored volumes.

**Exclude volumes in mirrored volume relationships**

Volumes with mirrored volumes are not analyzed and no recommendations are generated to move the volumes to other storage pools.

**Include volumes in mirrored volume relationships**

When a volume with a mirrored volume is moved to another pool, the primary volume is removed. The original secondary volume in its current pool is designated as the secondary volume. The copy of the original secondary volume in the destination pool is designated as the primary volume.

4. Specify the number of days that are used to estimate the activity levels of the pools. By default, the value for the activity level of pools is estimated based on the performance data that was collected on the previous day.
5. Specify the placement of optimized volumes in capacity pools. These options are only available if you select one or more volumes in capacity pools.
  - Yes** Restrict the placement of volumes that are optimized to storage pools that are in the same capacity pool as the volumes that are selected for analysis.
  - No** Restrict the placement of volumes that are optimized to the storage pools that were selected for analysis.
6. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool.

**Yes** If one or more of the volumes that are assigned to the same server or hypervisor require optimization, recommendations are only generated if all of the related volumes can be placed in the same destination pool.

**Multiple host connections to the same hypervisor or server:** If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

**No** If one or more of the volumes that are assigned to the same server or hypervisor require optimization, recommendations might be generated to place the related volumes in different destination pools.

7. Optional: To set activity limits for each tier, click **Advanced balancing**.
8. Click **Analyze**.

The Balance Pools wizard opens.

When you close the Balance Pools wizard, a task is created that shows the recommendations on the Balance Analysis page. You can implement the recommended changes, or you can create a schedule to implement the changes later.

**Analyzing volumes with service classes:** Recommendations are not generated to move volumes that are assigned a service class unless the destination pool can meet the requirements of the service class.

## Modifying the criteria for balancing pools

When a task is created after you run the Balance Pools wizard, you can modify the criteria that determine which pools are balanced.

You must run the Balance Pools wizard to create a balance analysis task.

The criteria that you modify are used to analyze the storage pools that you selected when you ran the Balance Pools wizard.

1. On the Navigation pane, click **Home > Tasks**.
2. Right-click a balance analysis task, and click **Edit**.
3. Make your changes, and then click **Analyze**.

Based on the criteria that you specified, the pools are analyzed to determine whether the pools require balancing. If one or more pools require balancing, recommendations are generated to redistribute the volumes to pools on the same tier of storage.

## Criteria for identifying the pools that require balancing

To identify the pools that require balancing, performance data is collected to determine the relative activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

To determine the relative activity of pools on the same tier and on the same back-end storage system, the following values are calculated:

- The current activity level of the pool is calculated.
- The activity deviation percentage of the pool is calculated.

**Tip:** On the Pools page, the activity level of the pool is shown in the **Activity** column. The difference between the activity level of the pool and all of the other pools on the same tier and storage system is shown in the **Activity Deviation (%)** column.

The following formulas are used to calculate the activity level and the activity deviation percentage for pools.

*Table 73. Formulas for calculating the activity level and activity deviation percentage values for pools*

Value calculated	Formula that is used in the calculation
Activity level of the pool	$[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$
Activity level of the pool on XIV systems	$(\text{Total I/O Rate} \div \text{Total Capacity})$

Table 73. Formulas for calculating the activity level and activity deviation percentage values for pools (continued)

Value calculated	Formula that is used in the calculation
Activity deviation percentage of the pool	$[(\text{Pool Activity Level} - \text{Average Activity Level}) \div \text{Average Activity Level}] \times 100$ <p>A positive value indicates that the relative activity level of the pool is above average. A negative value indicates that the relative activity level of the pool is below average.</p>

If the activity deviation percentage of the pool lies within the range of <-10% to >10%, the pool is considered to be balanced. If the activity deviation percentage of the pool exceeds the deviation threshold of 10%, the pool is a candidate for balancing.

For example, you want to identify the tier-1 pools that require balancing on a storage environment that consists of a Storwize V7000 Unified storage virtualizer with one DS8000 storage system. You open the Pools page and view the activity deviation percentage values for the tier-1 pools.

The Pools page shows that three tier-1 pools have the following activity levels and activity deviation percentages.

Table 74. Tier-1 pools activity and activity deviation percentages

Pool	Activity	Activity deviation percentage
pool_1	110	-3%
pool_2	90	-20%
pool_3	140	24%

The pool\_3 pool is a candidate for balancing because the activity deviation value for the pool is 14% higher than the deviation threshold of 10%.

The pool\_1 pool is considered to be balanced because the activity deviation percentage of the pool lies within the range for balanced pools. The pool\_1 pool and the pool\_2 pool, are potential target pools for the volumes that must be moved from the pool\_3 pool.

When you add volumes to pools, you want to ensure that the performance level of the pool continues to meet the workload requirements of the volumes. To maintain the performance capability of the pools, you must set an activity limit for each tier of storage.

To help you specify activity limits, the following information is provided for each tier of storage in the Balance Pools wizard, and on the Optimize the Placement of Volumes page of the Analyze Tiering wizard:

- The current average activity of all of the pools on the tier
- The current highest activity level

You can use the information that is provided to specify the activity limits for each tier. Alternatively, you can click **Help me calculate** and set the maximum threshold

rates for disk read and write I/O operations. With the information that you provide, the activity limit for the pool is calculated using the following formula:

$$(\text{Max Read I/O Rate} + \text{Max Write I/O Rate}) \div \text{Average Capacity of Pools}$$

---

## Transforming and migrating volumes

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Before you open the Transform Storage wizard, you must probe the storage virtualizers to collect information about the logical devices, such as volumes and virtual disks that are associated with the storage virtualizer.

You can convert or move volumes in pools on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

For volumes in mirrored volume relationships, you can convert the primary volume, or the secondary volume, or both volumes.

**Restrictions:** The restrictions for converting or moving volumes are as follows:

- To move volumes, the destination pool must be on the same storage virtualizer. You cannot move a volume from a pool on one storage virtualizer to a pool on another storage virtualizer.
- To convert or move volumes, image mode volumes must be converted to managed mode volumes.
- To convert fully allocated volumes to compressed volumes, you must have the IBM Real-time Compression™ license.

**Migrating volumes with service classes:** A warning is generated when you move one or more volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

1. In the menu bar in the web-based GUI, go to **Storage > Pools**.
2. Click two or more storage volumes, and then select **Transform Storage** from the context-sensitive menu. Alternatively, you can click **Actions > Transform Storage**.

The Transform Storage wizard opens.

When you close the Transform Storage wizard, the Transform Plan page is opened. A task is created which you can run to implement the recommendations.

**Related tasks:**

“Modifying the period for analyzing performance data” on page 281

To change the default period for analyzing performance data in storage virtualizer pools, issue the `setdscfg` command. In the Transform Storage wizard, the period that is set for analyzing performance data is used to calculate the average workload activity of the volumes that are selected for analysis.

---

## Modifying the period for analyzing performance data

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. In the Transform Storage wizard, the period that is set for analyzing performance data is used to calculate the average workload activity of the volumes that are selected for analysis.

You must have Administrator authority to use the **setdscfg** command.

When you run the Transform Storage wizard, the performance data that is collected about the pools on storage virtualizers is analyzed. By default, the period for analyzing performance data is set to 14 days. To provide a more accurate analysis of the workload activity of the volumes in your storage environment, you can extend or shorten the period.

To change the analysis period, complete the following steps:

1. To start the Tivoli Storage Productivity Center CLI, issue the **tpctool** command.
2. Issue the **setdscfg** command, and then set the following values by using the **-property** and the **-context** parameters:
  - a. Specify *tiering\_pm\_days* as the property value.
  - b. Specify *tiering* as the context value, and then specify the number of days for the analysis period.

For example, to change the analysis period from the default value of 14 days to 28 days, issue the following command:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****  
-property tiering_pm_days -context tiering 28
```

The default value is changed to the number of days that you specified, for example 28.

**Tip:** To verify that the analysis period is changed, you enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd ***** -context tiering
```

---

## Modifying the properties of resources

In the web-based GUI, you can modify the properties of the resources that are monitored by Tivoli Storage Productivity Center.


The following information summarizes the resources and properties that you can modify, and the steps for modifying each type of resource.

Resource	Properties that can be modified	Steps
Storage systems	<ul style="list-style-type: none"> <li>• Name (label)</li> <li>• Location</li> <li>• Custom tags (3)</li> </ul>	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, go to <b>Storage &gt; Storage Systems..</b></li> <li>2. Right-click the storage system that you want to modify and select <b>View Properties.</b></li> <li>3. On the properties notebook, click <b>Edit.</b></li> <li>4. Modify the properties and click <b>Save.</b></li> </ol>
Pools	<ul style="list-style-type: none"> <li>• Tier</li> <li>• Back-end storage system type</li> <li>• Back-end storage RAID level</li> <li>• Back-end storage disk type</li> <li>• Back-end storage disks</li> </ul> <p><b>Tip:</b> You can edit the back-end values if the back-end storage system for the related storage virtualizer was not probed and the values for the pool are unknown. If you manually enter values for these properties, Tivoli Storage Productivity Center uses the values to help calculate the approximate read I/O capability of the pool.</p>	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, go to <b>Storage &gt; Pools.</b></li> <li>2. Right-click the pool that you want to modify and select <b>View Properties.</b></li> <li>3. On the properties notebook, click <b>Edit.</b></li> <li>4. Modify the properties and click <b>Save.</b></li> </ol>



Resource	Properties that can be modified	Steps
Servers	<p><b>Servers with Storage Resource agents</b>            You can modify the following properties for servers where agents are deployed:</p> <ul style="list-style-type: none"> <li>• Location</li> <li>• Custom tags (3)</li> <li>• Fabric functions</li> <li>• Agent trace</li> <li>• Trace level</li> <li>• Number of trace files</li> <li>• Trace file size</li> <li>• Run scripts on agent</li> </ul> <p><b>Servers without Storage Resource agents (agentless)</b>            You can modify the following properties for agentless servers:</p> <ul style="list-style-type: none"> <li>• Name (label)</li> <li>• OS type</li> <li>• IP address</li> <li>• Virtual machine</li> <li>• Location</li> <li>• Custom tags (3)</li> </ul>	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, go to <b>Servers &gt; Servers</b>.</li> <li>2. Right-click the server that you want to modify and select <b>View Properties</b>.</li> <li>3. On the properties notebook, click <b>Edit</b>.</li> <li>4. Modify the properties and click <b>Save</b>.</li> </ol>
Hypervisors	<ul style="list-style-type: none"> <li>• Location</li> <li>• Custom tags (3)</li> </ul>	<ol style="list-style-type: none"> <li>1. In the menu bar in the web-based GUI, go to <b>Servers &gt; Hypervisors</b>.</li> <li>2. Right-click the hypervisor that you want to modify and select <b>View Properties</b>.</li> <li>3. On the properties notebook, click <b>Edit</b>.</li> <li>4. Modify the properties and click <b>Save</b>.</li> </ol>

Resource	Properties that can be modified	Steps
Switches	<ul style="list-style-type: none"> <li>Name (label)</li> <li>Location</li> <li>Custom tags (3)</li> </ul>	<ol style="list-style-type: none"> <li>In the menu bar in the web-based GUI, go to <b>Network &gt; Switches</b>.</li> <li>Right-click the switch that you want to modify and select <b>View Properties</b>.</li> <li>On the properties notebook, click <b>Edit</b>.</li> <li>Modify the properties and click <b>Save</b>.</li> </ol>
Fabrics	<ul style="list-style-type: none"> <li>Name (label)</li> <li>Location</li> <li>Custom tags (3)</li> </ul>	<ol style="list-style-type: none"> <li>In the menu bar in the web-based GUI, go to <b>Network &gt; Fabrics</b>.</li> <li>Right-click the fabric that you want to modify and select <b>View Properties</b>.</li> <li>On the properties notebook, click <b>Edit</b>.</li> <li>Modify the properties and click <b>Save</b>.</li> </ol>

**Tip:** To view the description of a property that you can modify for a resource, hover the mouse pointer over the hover icon  for that property in the web-based GUI.

---

## Opening the management GUI for a resource

Resources such as storage systems and switches can have their own management GUIs or element managers. In Tivoli Storage Productivity Center, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by Tivoli Storage Productivity Center.
- The resource has a management GUI that is web-based or can be accessed with Java Web Start.

- In the menu bar in the web-based GUI, go to a type of top-level resource. For example, if you want open the GUI for a storage system, go to **Storage > Storage Systems**. If you want open the GUI for a switch, go to **Network > Switches**.
- Right-click the resource and select the action to open the GUI. For storage systems, select **Open Storage System GUI**. For a switch, select **Open Switch GUI**. The start page for the management GUI is opened in a separate web browser window.

## Opening the DS Storage Manager GUI for a DS8000

If the firmware level for a DS8000 is earlier than R6.2, you must complete extra configuration steps before you can open its management GUI. To determine the firmware version of a DS8000 and enable the ability to open its management GUI, complete the following steps:

1. Start the Tivoli Storage Productivity Center web-based GUI.
2. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**.
3. Right-click the DS8000 and select **View Properties**.
4. Check the version in the **Firmware** field. If the firmware version is earlier than R6.2, continue to the next step. If the version is R6.2 or later, you do not need to complete this procedure.
5. Download and start the PuTTY utility. PuTTY is a free implementation of Telnet and SSH for Windows and UNIX. You can download PuTTY from <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>.
6. On the Session page in the PuTTY tool, complete the following actions:
  - a. In the **Host Name (or IP address field)**, enter the IP address of the Hardware Management Console (HMC) for the DS8000. You can determine this IP address by viewing the properties of the DS8000 in the Tivoli Storage Productivity Center web-based GUI.
  - b. Select **SSH** for a **Connection type**.
  - c. Click **Open**.
7. Enter the user name and password for logging in to the HMC.
8. Run the following command: `touch /opt/esshmc/data/config/ntvbrsen`
9. In the web-based GUI, right-click the DS8000 and select **Open Storage System GUI**.

---

## Exporting information about resources, jobs, and alerts

You can save information about resources, jobs, or alerts to a PDF, CSV, or HTML file. Information that you can export from the web-based GUI includes all the values that are being shown in the columns for a list of resources, jobs, or alerts.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the web-based GUI has a column or set of columns that determines the order of its rows.
- Filter the list to show only the resources, jobs, or alerts that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the following pages:

### Resource list and resource details pages

You can export information about the top-level resources, internal resources, or related resources that are shown on resource list pages and

resource details pages. For top-level resources, you can also export information about related jobs or alerts.

For example, you can export information about monitored switches on the Switches page. When you select the export action, the generated file includes the column values for each monitored switch in the list. In the **Jobs** or **Alerts** tabs on the Switches page, you can export information about all the jobs or alerts that are related to switches.

### The Alerts page

On the **Home > Alerts** page, you can export information about all the alerts that were detected for monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of alerts.

### The Jobs page

On the **Home > Jobs** page, you can export information about all the jobs that are associated with monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of jobs.

1. In the web-based GUI, go to a list of resources, jobs, or alerts for which you want to export information.
2. Click **Actions** and select **Export > file\_format**, where *file\_format* represents the format that you want to use for a file. The following formats are available:

#### CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name","Column_name","Column_name","Column_name","Column_name","Column_name"  
"data","data","data","data","data","data"
```

where `Column_name` represents the name of a column in a table and `data` represents the data that is associated with a column.

**PDF** You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

**Tip:** When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

#### HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

**Related concepts:**

“Viewing information about resources” on page 185

You can view detailed information about the resources that are monitored by Tivoli Storage Productivity Center. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

“Customizing lists of resources, tasks, and alerts”

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

---

## Customizing lists of resources, tasks, and alerts

The web-based GUI uses lists to show information about resources, tasks, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.


**Related reference:**

“How performance information is displayed” on page 210

The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

### Filtering lists

You can filter the items that are shown in a list. Use filtering when you want to limit the rows that are shown in a list based on values in a specific column or set of columns. For example, when viewing alerts, you can filter the list so that only alerts with a value of Warning in the Severity column are shown.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to filter.
2. Click the **Filter** icon  to determine how to filter the list. The following options are available:

#### **Filter by Default Columns**


Select this option to apply the filter to the set of the default columns in a list. Each list contains a set of default columns on which you can filter.

For example, in the lists of alerts on the **Home > Alerts** page, select this option to apply filter text to the following columns at the same time: Condition, Severity, Alert Category, Resource, Internal Resource, Total Occurrences, Alert Name.

#### **Filter by a Specific Column**

Select this option to apply filter text to a specific column in a list. For example, on the **Home > Alerts** page, select the **Condition** column if you want to apply filter text to values in that column only.

**Restriction:** Some columns in lists cannot be used for filtering and are not shown under the **Filter by a Specific Column** option.

3. In **Filter** text box , type the text for filtering a list. You can type letters, whole words, partial words, and numbers. When filtering on certain columns, you can select from a list of values rather than typing in filter text.

For example, when filtering on the **Status** column, you can select one of the following values as the filter text: Error, Unreachable, Warning, Normal, Unknown.

4. Press Enter to apply the filter to the list. Only rows with values that match or partially match the filter text are shown.

For example, on the **Home > Alerts** page, if you select to filter on the Severity column and type critical as the filter text, only alerts that have a severity of critical and critical - acknowledged are shown in the list.

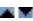
**Tip:** The name of a column to which a filter applies is shown next to the **Filter** icon .

5. Optional: To remove the filtering for a list, click **Reset**.

## Sorting lists

Each list of resources in the web-based GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.

For example, the list of ports on a fabric is ordered by switch name, blade slot, and port number, in that order. You can change the sort order of these columns to organize a list of resources according to your requirements.

**Tip:** A visual indicator  is shown next to the primary column that determines the sort order of a list.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to sort.
2. Complete the following tasks to sort the rows in the list:
  - To change whether the rows in a list are sorted by ascending or descending order, click the name of the primary column in the heading row. Click the name of the column again to reverse the sort order.
  - To change the column that determines the sort order of resources, click the name of that column. For example, click the Name column to sort the list of resources alphabetically, or click the **Status** column to sort the list according to the status of resources.

**Tip:** When you click a column other than the primary sort column, that column becomes the new primary sort column. However, the previous primary column maintains its sort order (ascending or descending).

For example, if a list of resources is ordered by the Name column in ascending order and you click the Status column, the rows are sorted according to the statuses of resources. At the same time, the resources for each type of status are still sorted in ascending order by the Name column. There is a maximum of three columns that can be sorted upon at the same time.

3. Optional: To reset the order of the rows for a column to the default setting, right-click the heading row in the list and select **Reset Table Preferences**.

## Showing, hiding, and reordering columns in lists

Each list in the web-based GUI has a set of columns that contain information about resources, jobs, or alerts. You can change the order of columns, show columns that are hidden, and hide columns that are shown in these lists.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to customize.
2. Complete the following tasks to customize the columns in the list:
  - To show or hide the columns in a list, right-click any of the column headings and select or clear the check box next to a column name.
  - To change the order of columns, click a column heading and drag it to a new position in the list.

**Tip:** The changes that you make to column settings are saved from session to session.

3. Optional: To reset the order and display of columns to the default setting, right-click any of the column headings and select **Restore Default View**.

---

## Managing tasks

Tasks are used to provision storage and optimize resources in your storage environment. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to provision storage and optimize resources. Use resource list and resource details pages to manage the tasks for specific resources and resource types.

Tasks are created when you take any of the following actions in the web-based GUI:

- Assign storage to servers or hypervisors by using the Provision Storage wizard.
- Move volumes to higher or lower tiers by using the Analyze Tiering wizard.
- Move or convert volumes by using the Transform Storage wizard.
- Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard.

Use the following actions to manage tasks:

- View the status of tasks in a central location.
- Identify all the tasks that are related to a specific resource or a type of resource.
- Identify all the tasks that are encountering problems when the tasks are running.
- Administer the tasks. The actions that are available depend on the type of task that you select. For example, you can run or rename a tiering-analysis task, or export provisioning task details to a text file.
- Change the number of task runs that are displayed.

**Tips:**

- In the stand-alone GUI, the term *schedules* or *jobs* is used instead of *tasks*.
- A *task run* is a single invocation of a task.


## Viewing tasks

Use the Tasks page in the web-based GUI to view all the tasks that are used by Tivoli Storage Productivity Center to provision storage and optimize resources. Use resource list and resource details pages to view the tasks for specific resources and resource types.

## Viewing all tasks

Use the Tasks page as a central location to view the tasks that Tivoli Storage Productivity Center is using to provision storage and optimize resources. You can view all the tasks for resources that are monitored by Tivoli Storage Productivity Center.

1. In the menu bar in the web-based GUI, go to **Home > Tasks**. The information about tasks is organized into columns. These columns include information about the status of tasks, the date and time of the last run for tasks, and the schedule for tasks.
2. Optional: View the status icons at the top of the page to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

**Tip:** To view descriptions of the icons and columns on the Tasks page, click the **Help** icon  in the upper-right corner of the page.

## Viewing pending tasks

Some tasks might be pending because they can be run or scheduled only by an administrator. If tasks are pending, the number of pending tasks is displayed in the title bar of the web-based GUI.

You must be assigned to the Administrator role to view the number of pending tasks in the title bar of the web-based GUI.

Pending tasks are tasks that were not run and are not scheduled to run. Pending tasks might be created by a user who is not authorized to run or schedule the task. For example, users who are assigned to the Monitor or External Application role can have permission to create provisioning tasks by using a service class. However, if the service class specifies that administrator approval is required, these users can only create the provisioning task. The provisioning task is pending until an administrator runs or schedules it.

To view pending tasks, complete the following steps:

1. In the title bar of the web-based GUI, look for the pending task count. If there are any pending tasks, the number of pending tasks is shown in the title bar of the web-based GUI. If no number appears in the title bar, there are no pending tasks.
2. To open a Tasks page that shows only the pending tasks, click the number of pending tasks in the title bar.

## Viewing tasks for a type of resource

Use resource list pages to view only the tasks for a specific type of top-level resource. Each resource list page has a **Tasks** tab that lists only the tasks for that type of resource. For example, you can use the list page for storage systems to view only the tasks that are associated with storage systems.

1. In the menu bar, select a type of top-level resource. For example, to view the tasks that are related to storage systems, go to **Storage > Storage Systems**.

**Exception:** Provisioning tasks that provision network-attached storage (NAS) file shares to servers or hypervisors are not shown on the resource list pages for servers or hypervisors. You must use the Tasks page to view these tasks.

2. Click the **Tasks** tab to view a list of all the tasks for the monitored resources of that type.




The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.

**Tip:** The status icon on the **Tasks** tab represents the most critical status of the tasks that are associated with the monitored resources. For example, if one of the tasks for storage systems failed, the following icon is shown on the tab:



### Tasks

- Optional: View the status icons on the **Tasks** page for a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

**Tip:** To view descriptions of the icons and columns on the **Tasks** page, click the **Help** icon .

## Viewing tasks for a specific resource


Use resource detail pages to view only the tasks for a specific monitored resource. Each resource detail page has a **Tasks** link that lists only the tasks for the associated storage system, server, hypervisor, switch, or fabric.

- In the menu bar, select a type of top-level resource. For example, if you want to view the tasks for a storage system, go to **Storage > Storage Systems**.
- Right-click a resource in the list and select **View Details**. The resource detail page is displayed. In the **General** section, the number next to **Tasks** represents the number of tasks that are associated with the resource. For example, **Tasks (3)** indicates that you can view three tasks.

The icon next to the **Tasks** link represents the most critical status of the tasks that are associated with the resource. For example, if a task failed, the following icon is shown:



- Click **Tasks** in the **General** section to view the tasks for the resource. The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.
- Optional: View the status icons on the Tasks pane to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.


**Tip:** To view descriptions of the icons and columns on the Tasks pane, click the **Help** icon .

## Viewing task details

Use task detail pages to view detailed information about the tasks that Tivoli Storage Productivity Center uses to optimize resources and provision storage. You can use this information to troubleshoot the reason why a task failed, analyze details of a task before you implement task recommendations, or schedule future task runs.

On all the task detail pages, you can view information such as the task status, the duration and completion date of task runs, and the actions that you can take to manage the task. For tasks such as tiering analysis, balance analysis, and transform plan, you can view information about the recommendations that are generated by

the most recent or earlier task runs. For provisioning tasks, you can view the status of the actions that the provisioning task completes for each of the resources that are configured in the task.




1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to storage systems, go to **Storage > Storage Systems**, and click the **Tasks** tab.
2. Locate the task that you want to view in more detail.
3. Right-click the task row and select **View Details**. The task details page opens in a separate window. For example, the Tiering Analysis page opens if you select a tiering-analysis task row.
4. Optional: To view task logs, click **Open Logs**. You can use the information in the log file to troubleshoot any errors that might occur when a task is run.
5. Optional: For provisioning tasks, you can export information about the provisioning task to a text file for offline viewing. To export information, click the  icon at the top of the Provisioning page.

## Viewing task logs

Task logs include detailed information about the status, actions, and progress of a task. Status information includes informational, warning, and error messages that are related to each action in a task that is taken during processing. You can use this information to troubleshoot any errors that might occur when a task is run.

The following examples show the type of information that you can view in the task logs:

- View the number of recommendations that are generated by tiering-analysis, balance-analysis, or transform-plan task runs.
- View the options that you select when you use wizards, such as the Transform Storage, Analyze Tiering, or Balance Pools wizards. For example, when you use the Analyze Tiering wizard, a tiering-analysis task is generated and the options that you choose in the wizard are recorded in the task log.
- View error messages that are related to task processing.

1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to storage systems, go to **Storage > Storage Systems**, and click the **Tasks** tab.
2. Locate the task that has a log file that you want to view.
3. Right-click the task row and select **Open Logs**. The Logs page shows log entries for the most recent task run.
4. Optional: On the Logs page, to view the log for a previous task run, select a task run from the **Select a log** list.
5. Optional: To view only the actions in a task run that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only actions that have the following statuses:
  -  Only error entries
  -  Only warning entries
  -  Error and warning entries
6. Optional: To view an explanation of the message that is associated with an action, click the link in the **ID** column.

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a task. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the task run.
- The status for each action in a task.
- The date and time when an action was completed. The date, time, and time zone of the action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

**Tip:** You can use the navigation and search functions of the web browser to locate information on the page.

### Setting the number of task runs that are displayed

You can change the number of task runs that are shown in the web-based GUI by configuring the settings on the History Retention page. By default, the information for the last five runs of a task is shown.

1. In the menu bar in the web-based GUI, go to **Settings > History Retention**.
2. Click **Edit**.
3. In the **Job logs** field, enter the number of runs that you want to retain for display.
4. Click **Save** to save your changes.

After you update the value for job logs, the list of task runs in the web-based GUI is automatically updated within a few minutes.

## Managing tasks for provisioning

Provisioning tasks are used to create storage volumes and assign the volumes to servers or hypervisors. Provisioning tasks are also used to create network-attached storage (NAS) file shares. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to provision storage. Use resource list and resource details pages to manage the provisioning tasks for specific resources and resource types.

Use the web-based GUI to complete the following provisioning actions:

- Create provisioning tasks by using the Provision Storage wizard
- Run provisioning tasks immediately
- Schedule provisioning tasks to run at a specified date and time
- Export information about provisioning tasks to text files
- Delete provisioning tasks

#### Related concepts:

“Provisioning storage” on page 75

The Provision Storage wizard simplifies the steps for submitting provisioning requests. When you use the Provision Storage wizard, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require.

## Creating provisioning tasks

Provisioning tasks are created when you use the Provision Storage wizard to assign storage volumes or NAS file shares to servers or hypervisors.

To use the Provision Storage wizard, you must have Administrator privileges.

### Related tasks:

“Provisioning storage with the Tivoli Storage Productivity Center GUI” on page 262

To provision volumes or shares to servers, hypervisors, and clusters, use the Provision Storage wizard.

“Collecting information about shares on storage systems” on page 159

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

## Running provisioning tasks

Use the Tasks page to run provisioning tasks. You can run provisioning tasks immediately or you can schedule tasks to run at a date and time that you specify. When you run a provisioning task, the configuration changes that are contained in the task are implemented. For example, storage volumes are created and assigned to servers.

To run provisioning tasks, you must be the owner of the task or have Administrator privileges.

Provisioning tasks can run only once. You can run the task when you complete the steps of the Provision Storage wizard or you can save the provisioning task and run the task later.

The **Run Now** and **Schedule** actions are available for a provisioning task when the task is not running, and was not previously run.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the provisioning tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the provisioning task that you want to run.
3. Depending on when you want the provisioning task to run, choose one of the following options:
  - To run the provisioning task immediately, right-click the task row and select **Run Now**. The task status changes to Running and the provisioning task runs immediately.
  - To schedule the provisioning task to run at a specific date and time, complete the following steps:
    - a. Right-click the task row and select **Schedule**.
    - b. From the Schedule page, select a future date and time from the **Provision** lists, and click **Save**.

You can also run or schedule the provisioning task from the Provisioning detail page. On the Provisioning detail page, click **Execute** to run the task immediately or click **Schedule** to schedule the task to run at a specific date and time.


The provisioning task runs immediately or on the scheduled date and time. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View detailed information about the progress of the task on the Provisioning detail page. You can view the task duration and the status of the actions that the provisioning task completes for each of the resources that are configured in the task.
- View the task logs from the Tasks page or from the Provisioning detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

### Exporting provisioning task information

Use the Provisioning detail page to export details of a provisioning task to a text file. The task details that are shown on the Provisioning page are saved to the generated text file.

If you do not have authority to provision storage, you can export the provisioning task details to a file and email the file to a user with Administrator privileges. The administrator can run or schedule the provisioning task if they approve the task details.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the provisioning tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the provisioning task that you want to export.
3. Right-click the task row and select **View Details**.
4. From the Provisioning page, to start the export, click the  icon in the upper-right corner of the page.
5. Save the generated text file. Depending on how your web browser is configured, you can specify the file name and the location where the file is saved.

The provisioning task information is exported to a text file. The generated file includes all the provisioning task details that are listed on the Provisioning page. Columns that are hidden on the Provisioning page are included in the generated file. The information in the file is organized according to a default order and is not affected by the column order that is defined on the Provisioning page.

## Managing tasks for tiering storage, balancing pools, and transforming storage

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

To schedule, run, edit, or delete tiering storage, balancing pools, and transforming-storage tasks, you must have Administrator privileges.

To complete any of the optimization tasks, such as balancing pools, you must complete the following steps:

1. Create the analysis task. Use a Tivoli Storage Productivity Center wizard to create the optimization analysis task. For example, use the Balance Pools wizard to create a balance-analysis task.
2. Run the analysis task to generate the optimization recommendations. For example, run the balance-analysis task to generate the recommendations to move volumes.
3. Edit analyze tiering and balancing pool tasks. You can modify the thresholds that were set in the Analyze Tiering and Balance Pools wizards.
4. Implement the recommendations to optimize storage. For example, implement the recommendations to balance pools by running the analysis-execution task.

Use the Tasks, resource list, and resource detail pages in the web-based GUI to take the following actions for optimization tasks:

- Run the optimization analysis immediately to generate optimization recommendations.
- Schedule the optimization analysis to run at a specified date and time.
- Implement optimization recommendations immediately.
- Schedule the implementation of optimization recommendations for a specified date and time.
- Pause, cancel, and resume the implementation of optimization recommendations.
- Delete optimization tasks.
- View task logs.

**Related tasks:**

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

“Modifying the criteria for balancing pools” on page 278

When a task is created after you run the Balance Pools wizard, you can modify the criteria that determine which pools are balanced.

## Creating analysis tasks

To create tasks for optimizing resources in your storage environment, use Tivoli Storage Productivity Center wizards.

To use the wizards, you must have Administrator privileges.

Storage optimization tasks are created when you take any of the following actions in the web-based GUI:

### Tier storage

Move volumes to higher or lower tiers by using the Analyze Tiering wizard. A tiering-analysis task is created.

### Balance pools

Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard. A balance-analysis task is created.

### Transform storage

Move or convert volumes by using the Transform Storage wizard. A transform-plan task is created.

**Related tasks:**

“Tiering volumes by I/O density, I/O rate, and file usage” on page 273

To tier volumes by setting minimum thresholds for I/O rate or I/O density, and file usage, use the Analyze Tiering wizard.

“Balancing pools” on page 276

To distribute the workload of volumes across pools on the same tier, use the Balance Pools wizard. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier.

“Transforming and migrating volumes” on page 280

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

## Running analysis tasks

To run tiering analysis, balance analysis, and transform-plan tasks, use the Tasks page. When you run the analysis for any of these optimization tasks, recommendations are generated to optimize the storage resources that are configured in the task. For example, when you run a tiering-analysis task, recommendations are generated to tier volumes based on the criteria that you set in tiering policies.

To run tiering analysis, balance analysis, and transform-plan tasks, you must have Administrator privileges.

You can run tasks immediately or schedule tasks to run at a future date and time.

**Tip:** The **Run Now** action is available when the task is enabled and the task is not running.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tiering-analysis tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to run.
3. Right-click the task row and select **Run Now**. The task status changes to Running.
4. Optional: To schedule future task runs for tiering analysis or balance-analysis tasks, right-click the task row and select **Schedule Analysis**. On the Schedule Analysis page, select the time that the task is run and how often the task is run.

**Restriction:** You cannot schedule future task runs for transform-plan tasks.

5. Optional: To enable or disable a tiering analysis or balance-analysis task, select **Enabled** or **Disabled** on the Schedule Analysis page. When the task is enabled, the task runs according to the defined schedule. If the task is disabled, you can modify the task schedule but the task cannot run at the scheduled time.

When the task run completes, you can view the recommendations that are generated on the task detail page. For example, you can view the recommendations that are generated for a tiering-analysis task on the Tiering Analysis page. To check the progress of the task run, you can take the following actions:

- View the task status on the Tasks page.
- View information about the progress of the task on the task detail page. You can view the duration of the task run, and the date and time that the task run completed.

- View the task logs from the Tasks page or from the task detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

**Related tasks:**

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

## Implementing recommendations to optimize storage

Use the Tasks page and task detail pages to implement optimization recommendations to tier storage, balance pools, and transform storage. You can implement recommendations immediately or you can schedule the implementation of the recommendations for a future date and time.

**Related tasks:**

“Modifying the criteria for analyzing tiering” on page 274

When a task is created after you run the Analyze Tiering wizard, you can modify the criteria for tiering the volumes that were selected for analysis.

### Implementing recommendations immediately:

To implement optimization recommendations to tier storage, balance pools, and transform storage, use the task detail pages. For example, use the Balance Analysis page to implement the recommendations to balance pools.

To implement storage optimization recommendations, you must have Administrator privileges.

Tivoli Storage Productivity Center uses analysis-execution tasks to implement optimization recommendations. When you implement recommendations for a tiering analysis, balance analysis, or transform-plan task, the following process takes place:

- An analysis-execution task is automatically created so that the recommendations can be implemented.
- The analysis-execution task runs immediately and the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

You can view analysis-execution tasks on the Tasks page. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

The **Execute** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
  - An analysis-execution task was not already run for the latest set of recommendations.
  - An analysis-execution task is not running or paused.
1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the analysis-execution tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
  2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to implement.



3. Right-click the task row and select **View Details**. For example, the Balance Analysis page opens if you select a balance-analysis task.
4. Click **Execute**. An analysis-execution task is created and immediately implemented.

The analysis-execution task runs immediately. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View the status of individual task recommendations on the Analysis Execution page.
- View the task logs from the Tasks page or from the Analysis Execution page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

#### **Scheduling the implementation of recommendations:**

To schedule the implementation of optimization recommendations to tier storage, balance pools, and transform storage, use the Tasks page.

To schedule the implementation of storage optimization recommendations, you must have Administrator privileges.

The **Schedule Execution** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task is not running or paused.
- An analysis-execution task was not already run for the latest set of analysis recommendations. For each run of the tiering analysis, balance analysis, or transform-plan task, you can complete only one run of the analysis-execution task.

When you define the schedule, an analysis-execution task is created for the implementation of the optimization recommendations.

**Tip:** The analysis-execution task implements the recommendations that were current when the analysis-execution task was defined. For example, you define a tiering-analysis task that runs daily. You review the recommendations every day and, on a Tuesday, you schedule an analysis-execution task to run on the following Friday. When the analysis-execution task runs on Friday, the recommendations that were current on Tuesday are implemented, even if more recent recommendations are available.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tiering analysis tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to schedule.
3. Right-click the task row and select **Schedule Execution**.
4. On the Schedule Execution page, select the date and time that the recommendations are implemented and click **Save**.

An analysis-execution task is created that you can monitor on the Tasks page. The name of the analysis-execution task is based on the name of the related tiering analysis, balance analysis, or transform-plan task. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

When the analysis-execution task runs on the scheduled date, the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

## Pausing, resuming, and canceling the implementation of recommendations

Tivoli Storage Productivity Center uses analysis-execution tasks to implement optimization recommendations to tier storage, balance pools, and transform storage. To pause, resume, and cancel analysis-execution tasks, use the Tasks page.

To pause, resume, or cancel analysis-execution tasks, you must have Administrator privileges.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the analysis-execution tasks that are related to servers, click **Servers > Servers**, and click the **Tasks** tab.
2. Locate the analysis-execution task that you want to pause, resume, or cancel.
3. Right-click the task row and select **Pause**, **Resume**, or **Cancel**.

The **Pause** action is only available when the analysis-execution task is running.

The **Resume** action is only available when the analysis-execution task is paused.

The **Cancel** action is available when the task is running or paused.

When you pause, resume, or cancel analysis-execution tasks, the following processing occurs:

**Pause** The status of the analysis-execution task changes to Paused and the recommendations that are being implemented are suspended.

### Resume

The status of the analysis-execution task changes to Running and the recommendations are resumed.

### Cancel

The status of the analysis-execution task changes to Canceled and the recommendations that are being implemented are stopped. Any recommendations that are already implemented for this analysis-execution task are not rolled back. For example, a recommendation is generated to move a volume from pool\_a to pool\_b. If this recommendation is already implemented before the task is canceled, the volume remains in pool\_b. Use the **Analysis Execution** page to determine which recommendations were implemented for the task before the task was canceled.

You cannot restart an analysis-execution task that is canceled. If you want to implement the same set of recommendations again, you must complete the following steps:

1. Delete the analysis-execution task that you canceled.
2. Use the **Schedule Execution** action for the tiering analysis, balance analysis, or transform-plan task to create an analysis-execution task.

**Restriction:** The implementation of recommendations for optimizing the volumes is canceled except for the recommendations to optimize the volumes with mirrored volume copies. The implementation of recommendations for optimizing volumes with mirrored volume copies cannot be canceled.

## Renaming tasks

You can change the name of a task by using the task detail page. For example, use the Provisioning page to change the name of a provisioning task. A task name is automatically generated when the task is created. If the original name is ambiguous or is not clear enough, you can change the task name to a unique task name that you provide.

You must have the following privileges to rename the different task types:

### Provisioning

To rename provisioning tasks, you must be the owner of the task or have Administrator privileges.

### Other task types

To rename all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the task that you want to rename.
3. Right-click the task row and select **View Details**. The task detail page opens in a separate window. For example, the Provisioning page opens if you select a provisioning task.
4. Type the new name in the task name field. The naming conventions for task names are:
  - The maximum length is 64 characters.
  - Uppercase and lowercase letters are allowed.
  - Spaces and numbers are allowed.
  - The following special characters are allowed:  
! # % & \* + - / = ? ^ \_ { } ( )
5. Click **Rename**.

The task is renamed. You can view the renamed task on the Tasks page.

## Deleting tasks

Use the Tasks page to delete tasks, such as provisioning, tiering analysis, balance analysis, transform plan, and analysis-execution tasks.

You must have the following privileges to delete the different task types:

### Provisioning

To delete provisioning tasks, you must be the owner of the task or have Administrator privileges.

### Other task types

To delete all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

The following conditions apply for the different task types:

### Analysis-execution

You can delete analysis-execution tasks that were canceled, already run, or scheduled to run in the future. You cannot delete tasks that are running or paused.

### Other task types

You can delete all other tasks if the tasks were already run or if the tasks are scheduled to run in the future. For example, you can delete provisioning, tiering analysis, balance analysis, and transform-plan tasks that were already run. You cannot delete tasks that are running.

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view tasks that are related to servers, go to **Servers > Servers**, and click the **Tasks** tab.
2. Locate the task that you want to delete.
3. Right-click the task row and select **Delete**. Alternatively, you can delete the task by clicking **Delete** on the task detail page. The Delete page opens in a separate window.
4. If you are deleting a provisioning task that was not yet run, enter the reason that you are deleting the task.
5. Click **Delete**. To exit the window without deleting the task, click **Cancel**.

---

## Chapter 6. Managing replication

Manage replication components, including setting up the management servers, adding storage connections and paths, creating sessions and copy sets, setting up data replication, performing planned maintenance, and recovering from a disaster.

---

### Administering

Administer IBM Tivoli Storage Productivity Center for Replication to authorize users, start and use the graphical user interface (GUI), start and stop services, and many other administrative tasks.

### Starting and stopping Tivoli Storage Productivity Center for Replication

Use these procedures to start and stop Tivoli Storage Productivity Center for Replication, including the embedded WebSphere Application Server and DB2.

#### Starting the Tivoli Storage Productivity Center for Replication server

The commands that are required to start the Tivoli Storage Productivity Center for Replication server depend on the operating system.

The Tivoli Storage Productivity Center for Replication server is started by default after the application is installed.

#### Starting Tivoli Storage Productivity Center for Replication on Windows:

To start Tivoli Storage Productivity Center for Replication, issue the following command from the command prompt:

```
TPC_install_directory\scripts\startTPCReplication.bat
```

Where the variable *TPC\_install\_directory* is the installation directory. The default directory is C:\Program Files\IBM\TPC.

#### Starting Tivoli Storage Productivity Center for Replication on AIX or Linux:

To start Tivoli Storage Productivity Center for Replication, issue the following command from the command line:

```
/TPC_install_directory/scripts/startTPCReplication.sh
```

Where *TPC\_install\_directory* is the installation directory. The default directory is */opt/IBM/TPC*.

#### Starting Tivoli Storage Productivity Center for Replication on z/OS:

To start Tivoli Storage Productivity Center for Replication for System z, complete the following steps:

1. Ensure that the Tivoli Storage Productivity Center for Replication for System z HFS is mounted on the UNIX System Services. The root production directory is *path\_prefix/opt/Tivoli/RM*.
2. Issue the command from the IBM z/OS System Display and Search Facility (SDSF) panel to start Tivoli Storage Productivity Center for Replication for System z, as shown in the following example. The address spaces that are shown might be different in your environment.

```
/START IWNSRV
```

The initiator is IWNSRV and the job name is IWNSRV. For information about the IWNSRV job, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

## Stopping Tivoli Storage Productivity Center for Replication server

The commands that are required to stop the Tivoli Storage Productivity Center for Replication server depend on the operating system.

### Stopping Tivoli Storage Productivity Center for Replication on Windows:

To stop the Tivoli Storage Productivity Center for Replication server, issue the following command from the command prompt:

```
TPC_install_directory\scripts\stopTPCReplication.bat
```

Where the variable *TPC\_install\_directory* is the installation directory. The default directory is C:\Program Files\IBM\TPC.

### Stopping Tivoli Storage Productivity Center for Replication on AIX or Linux:

To stop Tivoli Storage Productivity Center for Replication, issue the following command from the command line:

```
/TPC_install_directory/scripts/stopTPCReplication.sh
```

Where *TPC\_install\_directory* is the installation directory. The default directory is /opt/IBM/TPC.

### Stopping Tivoli Storage Productivity Center for Replication on z/OS:

To stop Tivoli Storage Productivity Center for Replication for System z, complete the following steps:

1. Issue a command from the IBM z/OS System Display and Search Facility (SDSF) panel to stop Tivoli Storage Productivity Center for Replication for System z, as shown in the following example command. The address space that is shown might be different in your environment.

```
/STOP IWNSRV
```

Where */STOP IWNSRV* stops the server address space.

2. Issue the */D A,L* command to show the processes that are running and verify that the stop command completed. If the IWNSRV address space is displayed, issue the */CANCEL* command for the address space.

## Verifying that the Tivoli Storage Productivity Center for Replication server is running

You can verify that the Tivoli Storage Productivity Center for Replication server is running in several ways.

Complete one or more of these tasks to determine whether the Tivoli Storage Productivity Center for Replication server is running:

- Start the Tivoli Storage Productivity Center for Replication GUI or command-line interface shell. If either of these methods are successful, then the Tivoli Storage Productivity Center for Replication server is running.
- View the *csmTrace.log* file in the *TPC\_install\_directory\wlp\usr\servers\replicationServer\logs\CSM\* directory. If the *csmTrace.log* file is updated regularly and is increasing in size, then Tivoli Storage Productivity Center for Replication is running.

## Identifying the version of Tivoli Storage Productivity Center for Replication

The Tivoli Storage Productivity Center for Replication version determines the features that are available in the product.

To view the version of Tivoli Storage Productivity Center for Replication from the GUI, hover over the ? icon in the menu bar, and click **About**.

To view the version of Tivoli Storage Productivity Center for Replication from the command-line interface, use the `-ver` command.

## Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data

You can back up the entire IBM Tivoli Storage Productivity Center for Replication database. You can also back up copy sets in a specific session. You can then use the backup files to restore a previous configuration or recover from a disaster.

### Back up and recovery

You can back up and recover copy set data for a specific session and the complete Tivoli Storage Productivity Center for Replication database.

### Copy sets

You can export data about all copy sets in a specific session to maintain a backup copy that you can use to recover if you lose the session or to upgrade to a different management server.

When you export copy sets, Tivoli Storage Productivity Center for Replication creates a backup of the session and saves the data in a comma-separated value (CSV) file, which you can view or edit in a spreadsheet program such as Microsoft Excel. The exported CSV file includes the session name, session type, date that the data was exported, and the copy sets for each role pair. There is one line per copy set, and the volumes in the copy set are separated by a comma (for example: `ESS:2105.FCA57:VOL:17C7,ESS:2105.12043:VOL:17C7`).

The following example illustrates the content of the CSV file for a FlashCopy session. The first valid row must contain the appropriate role names for the session. The order of the copy sets does not matter, and you can include extra roles. A copy set is created from each row that follows the role names. All rows must have data in each column to be a valid row. The number sign (#) indicates that the line is a comment. Lines that are comments are ignored.

```
#Session1,  
#FlashCopy,  
#Oct 2, 2009 10:03:18 AM  
  
H1,T1  
DS8000:2107.FRLL1:VOL:1004,DS8000:2107.FRLL1:VOL:1104  
DS8000:2107.FRLL1:VOL:1011,DS8000:2107.FRLL1:VOL:1101  
DS8000:2107.FRLL1:VOL:1005,DS8000:2107.FRLL1:VOL:1105
```

**Important:** You must manually save this file on the local system when you export copy sets from the Tivoli Storage Productivity Center for Replication web interface.

## IBM Tivoli Storage Productivity Center for Replication database

Tivoli Storage Productivity Center for Replication database contains all product data, including data about storage systems, sessions, copy sets, paths, user administration, and management servers. You can back up this data and use the backup file to recover from a disaster or restore a previous configuration.

**Important:** You must have Administrator privileges to back up and recover the database.

The current data is stored in a new file each time you create a backup. The backup file is named *yyyyMMdd\_HHmmsSSS.zip*, where *yyyy* is the year, *MM* is the month, *dd* is the day, *HH* is the hour, *mm* is the minute, *ss* is the seconds, *SSS* is the milliseconds when the backup command was run. It is your responsibility to delete backup versions that are no longer needed.

By default, the backup file is stored in the following location:

### Tivoli Storage Productivity Center for Replication

```
TPC_install_directory\wlp\usr\servers\replicationServer\database\  
backup\
```

### Tivoli Storage Productivity Center for Replication for System z

```
path_prefix/opt/Tivoli/RM/wlp/usr/servers/replicationServer/  
database/backup/
```

You can change the default location by editing the **db.backup.location** property in the `rmsserver.properties` file, which is in the following location:

### Tivoli Storage Productivity Center for Replication

```
TPC_install_directory\wlp\usr\servers\replicationServer\properties\  

```

### Tivoli Storage Productivity Center for Replication for System z

```
path_prefix/opt/Tivoli/RM/wlp/usr/servers/replicationServer/  
properties/
```

The backup file contains the Tivoli Storage Productivity Center for Replication database data at the time the backup was completed. Any changes that were made after the backup are not reflected when the backup files are used to restore an Tivoli Storage Productivity Center for Replication database.

Create a new backup file in the following situations:

- After you change the Tivoli Storage Productivity Center for Replication database data, such as adding or deleting a storage system, changing properties, and changing user privileges
- After an Tivoli Storage Productivity Center for Replication session changes direction. For example, if a Metro Mirror session was copying data from H1 to H2 when the backup was taken, and later, the session was started in the H2 to H1 direction. The session must be in the Prepared state before you create the backup.
- After a site switch has been declared and the Enable Copy To Site or Set Production to Site command is issued. After you create a backup, consider deleting the previous backup to prevent Tivoli Storage Productivity Center for Replication from starting the copy in the wrong direction.

When you create a backup, ensure that all Tivoli Storage Productivity Center for Replication sessions are either in the Defined, Prepared, or Target Available state.



Restoring the Tivoli Storage Productivity Center for Replication database from a backup copy puts Tivoli Storage Productivity Center for Replication back to the point in time when the backup was made. Relationships that exist on the storage systems that were created by Tivoli Storage Productivity Center for Replication after the backup was made are no longer be managed by Tivoli Storage Productivity Center for Replication until you add the copy set to the session and Tivoli Storage Productivity Center for Replication assimilates the relationship into the session. Copy sets that were deleted after the backup are restored and a subsequent Start command to the session creates new relationships. Therefore, you must remove the deprecated copy sets before you issue the Start command.

After you restore a Global Mirror session, you must stop the Global Mirror master and subordinates before you restart the Global Mirror session. Refer to your storage system documentation for more information.

### Exporting copy set data

You can export data about all copy sets in a specific session to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Complete these steps to export the copy sets in a specific session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to export copy sets.
3. From the **Session Actions** list, select **Export > Export Copy Sets**. The Export Copy Sets window displays the status of the export and a link to the exported file if the export completed.
4. Right-click the link and save the file to a local system.

### Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to import copy sets.
3. From the **Session Actions** list, select **View/Modify > Add Copy Sets**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Click **Browse**, select the CSV file, and then click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets that you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and then click **Next**.
9. View the information on the Results page of the wizard, and then click **Finish**.

## Verifying the Java heap size

The Java heap size determines the number of volume pairs that Tivoli Storage Productivity Center for Replication can manage.

## Verifying the Java heap size for Tivoli Storage Productivity Center for Replication

The default Java heap size is 1024 MB, which supports fewer than 75,000 pairs. If you increase the Java heap size by 256 MB, support increases by 25,000 pairs. You can continue to increase the Java heap size until you reach the maximum heap size.

For more information about the maximum heap size, see the information about sizing the Java heap size in the IBM SDK for Java Tools information center at <http://publib.boulder.ibm.com/infocenter/javasdk/tools/index.jsp>.

For more information about how to set up the Java heap size, see the *WebSphere Application Server V8.5 Liberty profile* technote at <http://www.ibm.com/support/docview.wss?uid=swg21596474>.

## Verifying the Java heap size for Tivoli Storage Productivity Center for Replication for System z

The default Java heap size is 768 MB, which supports fewer than 50,000 pairs. If you increase the Java heap size by 256 MB, support increases by 25,000 pairs. You can continue to increase the Java heap size until you reach the maximum heap size.

For more information about the maximum heap size, see the information about sizing the Java heap size in the IBM SDK for Java Tools information center at <http://publib.boulder.ibm.com/infocenter/javasdk/tools/index.jsp>.

For more information about how to set up the Java heap size, see the *WebSphere Application Server V8.5 Liberty profile* technote at <http://www.ibm.com/support/docview.wss?uid=swg21596474>.

---

## Managing management servers

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

### Management servers

The *management server* is a system that has IBM Tivoli Storage Productivity Center for Replication installed. The management server provides a central point of control for managing data replication.

You can create a high-availability environment by setting up a *standby management server*. A standby management server is a second instance of Tivoli Storage Productivity Center for Replication that runs on a different physical system, but is continuously synchronized with the primary (or active) Tivoli Storage Productivity Center for Replication server. The active management server issues commands and processes events, while the standby management server records the changes to the active server. As a result, the standby management server contains identical data to the active management server and can take over and run the environment without any loss of data. If the active management server fails, you can issue the Takeover command to make the standby management server take over.

## Connecting the active management server to the standby management server

Ensure that the active management server is connected to the standby management server. This connection creates the *management server relationship* that begins the synchronization process. Each management server can be in only one management server relationship.

Problems might occur during the synchronization process. For example, the management server relationship might become disconnected for a number of reasons, including a connectivity problem or a problem with the alternate server.

If a problem occurs during the synchronization process, the alternate server database is restored to its original state before the synchronization process began. Issue the Reconnect command to restore synchronization.

## Completing a takeover on the standby management server

If you must complete a takeover and use the standby server, ensure that you shut down the active management server first. You must ensure that you do not have two active management servers. If there are two active management servers and a condition occurs on the storage systems, both management servers respond to the same conditions, which might lead to unexpected behavior.

If you complete an action on the active management server when the servers are disconnected, the servers will be out of synchrony.

## Viewing the status of the management servers

You can view the status of the active and standby management servers from the Management Servers panel in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). If you are logged on to the active management server, the icons on this panel show the status of the standby management server. If you are logged on to the standby management server, the icons on this panel show the status of the active management server.

When the status is Synchronized, the standby management server contains the same data that the active management server contains. Any update to the active management server database is replicated to the standby server database.

## Managing volumes on storage systems

When you add direct connections, Hardware Management Console (HMC) connections, or z/OS connections on the active management server, Tivoli Storage Productivity Center for Replication automatically enables the management of attached extended count key data (ECKD™) volumes, non-attached count key data (CKD) volumes, and all fixed-block volumes on the storage system. To disable management of volumes on the storage system, use the volume protection function.

## Setting exceptions for ephemeral port numbers

For management servers that use an Internet Protocol version 4 (IPv4) connection, the port that is shown on the Management Servers page of the Tivoli Storage Productivity Center for Replication GUI is the *listener* port. This port was set when Tivoli Storage Productivity Center for Replication was installed.

However, when the management server initiates a connection to standby server, the management server uses an *ephemeral* port.

An ephemeral port is a temporary port number that is assigned for the duration the connection. When the connection terminates, the ephemeral port is available for reuse. When the management server reconnects to the standby server, a new ephemeral port number is assigned.

If you are using a firewall, you must create exceptions that enable the management server to connect to the standby server by using outbound ephemeral ports.

## Information specific to management servers in z/OS environments

If the standby management server is not in the active server z/OS sysplex, the standby server is not able to communicate with the storage systems by using a z/OS connection. Therefore, another connection must be made by using a TCP/IP connection.

If DB2 is configured for data sharing mode across the z/OS sysplex, one of the Tivoli Storage Productivity Center for Replication servers must be configured to use the zero-administration embedded repository. If the embedded repository is not used, the two servers overwrite the same data in the Tivoli Storage Productivity Center for Replication database.

Tivoli Storage Productivity Center for Replication Basic Edition for System z does not support active and standby management servers. If you are using this product, the menu item **Management Servers** is not displayed on the **Settings** menu of the GUI.

## SNMP alerts

This topic describes the SNMP alerts that are sent by IBM Tivoli Storage Productivity Center for Replication and the associated object IDs (OIDs).

SNMP alerts are sent during the following general events:

- Session state change
- Configuration change
- Suspending-event notification
- Communication failure
- Management Server state change

### Session state change SNMP trap descriptions

This topic lists the SNMP traps that are sent during a session state change. A different trap is sent for each state change. These alerts are sent by only the active management server.

A session state change SNMP trap is sent each time the session changes to one of the following states:

- Defined
- Preparing
- Prepared
- Suspended
- Recovering
- Flashing
- Target Available
- Suspending

- (Metro Global Mirror only) SuspendedH2H3
- (Metro Global Mirror only) SuspendedH1H3

An SNMP trap is also sent when a recovery point objective (RPO) threshold is exceeded for a role pair that is in the session.

*Table 75. Session state change traps*

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.1	The state of session X has transitioned to Defined.
1.3.6.1.4.1.2.6.208.0.2	The state of session X has transitioned to Preparing.
1.3.6.1.4.1.2.6.208.0.3	The state of session X has transitioned to Prepared.
1.3.6.1.4.1.2.6.208.0.4	The state of session X has transitioned to Suspended.
1.3.6.1.4.1.2.6.208.0.5	The state of session X has transitioned to Recovering.
1.3.6.1.4.1.2.6.208.0.6	The state of session X has transitioned to Target Available.
1.3.6.1.4.1.2.6.208.0.19	The state of session X has transitioned to Suspending.
1.3.6.1.4.1.2.6.208.0.20	The state of session X has transitioned to SuspendedH2H3.
1.3.6.1.4.1.2.6.208.0.21	The state of session X has transitioned to SuspendedH1H3.
1.3.6.1.4.1.2.6.208.0.22	The state of session X has transitioned to Flashing.
1.3.6.1.4.1.2.6.208.0.23	The state of session X has transitioned to Terminating.
1.3.6.1.4.1.2.6.208.0.26	The recovery point objective for the role pair of X in session Y has passed the warning threshold of Z seconds.
1.3.6.1.4.1.2.6.208.0.27	The recovery point objective for the role pair of X in session Y has passed the severe threshold of Z seconds.

**Related reference:**

“Session states” on page 442

You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

**Configuration change SNMP trap descriptions**

This topic lists the SNMP traps that are sent when the configuration changes. These alerts are sent by only the active management server.

Configuration change SNMP traps are sent after the following configurations changes are made:

- One or more copy sets have been added or deleted from a session
  - An alert is sent for each set of copy sets added to or removed from a session. Note that an alert for copy set changes is sent only once within 15 minutes of a configuration change, so you might not see alerts from successive changes that occur within that 15-minute period. For example, if you make a copy set configuration change that causes an alert to be sent at 10:41:01, and if you were to make additional copy set changes at 10:42:04 and 10:50:09, no alerts would be sent for these two changes because they occurred within the 15-minute minimum interval from the first alert.
- PPRC path definitions have been changed
  - An alert is sent for each path configuration change made.

Table 76. Configuration change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.7	One or more copy sets have been added or deleted from this session. <b>Note:</b> An event is sent for each session at least every 15 minutes.
1.3.6.1.4.1.2.6.208.0.8	Peer-to-Peer Remote Copy (PPRC) path definitions have been changed. An event is sent for each path configuration change.

### Suspending-event notification SNMP trap descriptions

These SNMP traps that are sent during a suspending-event notification. The traps are sent by the active and standby management server.

Suspending-event notification SNMP traps indicate that a session has transitioned to a Severe status due to an unexpected error.

Table 77. Suspending-event notification traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.9	The session is in a Severe state due to an unexpected error.

### Communication-failure SNMP trap descriptions

This topic lists the SNMP traps that are sent during a communication-failure. These alerts are sent by both the active and standby management servers.

Communication-failure SNMP traps are sent after the following events occur:

- A server times out attempting to communicate with a storage system.
- A server encounters errors attempting to communicate with a storage system.
- An active server terminates communication with a standby server as a result of communication errors.
- A standby encounters communication errors with an active server.

After an SNMP trap for a given failure is sent, it is not resent unless communication has been reestablished and failed again.

Table 78. Communication-failure traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.10	Server X has timed out attempting to communicate with storage system Y.
1.3.6.1.4.1.2.6.208.0.11	Server X has encountered errors attempting to communicate with storage system Y.
1.3.6.1.4.1.2.6.208.0.12	Active server X has terminated communication with standby server Y as a result of communication errors.
1.3.6.1.4.1.2.6.208.0.13	Standby server X has encountered communication errors with active server Y.

### Management Servers state-change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the state of the management server changes. These alerts are sent by both the active and standby management servers.

A management server state change SNMP trap is sent each time the management server changes to one of the following states:

- Unknown
- Synchronization Pending
- Synchronized
- Disconnected Consistent
- Disconnected

Table 79. Management Servers state-change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.14	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Unknown (previously Offline).
1.3.6.1.4.1.2.6.208.0.15	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Synchronized.
1.3.6.1.4.1.2.6.208.0.16	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected Consistent (previously Consistent Offline).
1.3.6.1.4.1.2.6.208.0.17	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Synchronization Pending.
1.3.6.1.4.1.2.6.208.0.18	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected.

## Setting up a standby management server

You can set up a standby management server in two ways: setting up the management server you are logged in to as the standby management server, or designating another server as the standby management server.

**Note:** When you define a standby management server, the IBM Tivoli Storage Productivity Center for Replication code must be at the same level on both the standby and active management servers.

### Setting the local management server as the standby server

You can define the server on which you are currently logged in as the standby management server.

**Attention:** When you set a management server as the standby server, all the information on that management server is cleared. The operation cannot be undone.

**Important:** The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

To set the local management server as the standby server, complete the following steps:

1. In the menu bar, click **Settings > Management Servers**.
2. On the Management Servers page, from the **Select Action** list, select **Set this Server as Standby**.

3. Enter the domain name or IP address of the server that you want to use as the active management server.
4. Click **OK** to connect to the active server. The server on which you are logged in is now the standby server.

### Setting a remote management server as the standby server

You can define a remote server as the standby management server.

**Attention:** When you set a management server as the standby server, all the information on that management server is cleared. The operation cannot be undone.

**Important:** The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

To set a remote management server as the standby server, complete the following steps:

1. In the menu bar, click **Settings > Management Servers**.
2. On the Management Servers page, from the **Select Action** list, select **Define Standby**.
3. Enter the domain name or IP address of the server that you want to use as the standby management server. Log in to the standby management server by entering the user name and password.
4. Click **OK** to connect to the standby management server.

### Applying maintenance to an active management server

If an active management server develops a problem during an active session, you can reinstall Tivoli Storage Productivity Center for Replication on the server without affecting the session.

Follow these steps to reinstall Tivoli Storage Productivity Center for Replication on the active management server.

1. If the heartbeat is enabled, disable it:
  - a. In the menu bar, click **Settings > Advanced Tools**.
  - b. On the Advanced Tools page, click **Disable Heartbeat**.
2. On the standby server, Server 2, issue a take-over. This makes Server 2 the active server. It is possible that the original active server, Server 1, is still listed on the Server 2 Management Servers page. If so, select **Remove Standby**.
3. Disable the heartbeat on Server 2, in case there are any problems.
4. Uninstall Tivoli Storage Productivity Center for Replication on Server 1.
5. Reinstall Tivoli Storage Productivity Center for Replication on Server 1.

**Note:** If no changes have been made to the configuration while Server 1 was being reinstalled, steps 6 and 7 are not necessary.

6. When Tivoli Storage Productivity Center for Replication is running on Server 1, log into Server 2 and set Server 1 as the standby server for Server 2.

This step copies the configuration from Server 2 to Server 1. This process takes a few minutes.



7. When the management servers status is synchronized, issue a take-over on Server 1.

This makes Server 1 an active server, able to control sessions.

**Note:** It is possible that Server 2 is still listed on the Server 1 Management Servers page. If this is the case, select **Remove Standby**.

8. Disable the heartbeat on Server 1 to make sure this active server does not have any problems.

**Note:** If you do not need to reinstall Tivoli Storage Productivity Center for Replication on Server 2, skip steps 9 and 10.

9. Uninstall Tivoli Storage Productivity Center for Replication from Server 2.
10. Reinstall Tivoli Storage Productivity Center for Replication on Server 2.
11. On Server 2, go to the Management Servers page and select the **Set This Server As Standby** option, entering the information for Server 1. When this step is complete, Server 1 is the active server, and Server 2 is the standby server.
12. When you are confident that the active server is running without any problems, enable the heartbeat again, if desired.

## Reconnecting the active and standby management servers

If the active and standby management servers become disconnected, reestablish that connection.

To reconnect the active and standby management servers, complete the following steps:

1. In the menu bar, click **Settings > Management Servers**.
2. On the Management Servers page, from the **Select Action** list, select **Reconnect**.

## Performing a takeover on the standby management server

If the active management server fails, you can force the standby management server to take over monitoring and managing replication responsibilities.

**Important:** If the current active management server is still active, you must not attempt to control the replication environment simultaneously from both management servers. Instead, either reconfigure the current active management server to be a standby management server, or shut it down.

Perform these steps to cause the standby management server to become the active management server:

1. If the active management server is functioning, take it offline so you do not have two active management servers managing the same sessions.
2. Log on to the Tivoli Storage Productivity Center for Replication GUI that is running on the standby management server.
3. In the menu bar, click **Settings > Management Servers**.
4. On the Management Servers page, from the **Select Action** list, select **Takeover**.
5. To reestablish high-availability, complete one of these steps:
  - Choose another server to be the standby management server. See instructions for setting up a standby management server.

- Bring the failed management server back online, and then make that server the standby management server. See “Setting up a standby management server” on page 313.
- Bring the failed management server back online, and then make that server the active management server to return to the original configuration. Repeat the steps in this section and then add the original standby server as the standby server.

**Important:** Do *not* use the **Reconnect** command if you perform a takeover. You would use the **Reconnect** command when the active server loses its connection with the standby server; it reconnects the two servers. Do *not* use the **Reconnect** command after a takeover to reconnect to the original active server.

## Configuring SNMP

The SNMP community name has a default value of public.

To change the community name, modify or add the **csm.server.snmp\_community\_string** property in the `rmserver.properties` file. This file is in the `TPC_install_directory\wlp\usr\servers\replicationServer\properties\` directory.

## Adding SNMP managers

Use the **mksnmp** command to add an SNMP manager to the list of servers to which Tivoli Storage Productivity Center for Replication sends SNMP alerts.

Tivoli Storage Productivity Center for Replication uses management information base (MIB) files to provide a textual description of each SNMP alert that is sent by Tivoli Storage Productivity Center for Replication. You must configure the SNMP manager to use both the `SYSAPPL-MIB.mib` and `ibm-TPC-Replication.mib` files. These files are on the installation DVD in the `replication\CSM-Client\etc` directory. Follow the directions that are provided by your SNMP manager application to configure it to use the MIB files.

Tivoli Storage Productivity Center for Replication sends all SNMP alerts to each registered SNMP manager. SNMP alerts are not specific to any particular session, and all alerts for any session are sent. You cannot choose to send a subset of SNMP alerts.

## Changing the standby management server port number

The standby management server port is used for communication between the active and standby management server. This port is initially defined during the installation. You can manually change this port after installation.

**Important:** The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

1. Open the `rmserver.properties` file in the `TPC_install_directory\wlp\usr\servers\replicationServer\properties\` directory.
2. Modify the port number for the following property:  
`communications.haPort=port_number`

3. Restart Tivoli Storage Productivity Center for Replication. You must restart Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the Tivoli Storage Productivity Center for Replication management servers and must be done on each Tivoli Storage Productivity Center for Replication management server.

## Changing the client port number

The client port is used to log in to the graphical user interface and command-line interface from a remote system. This port is initially defined during the installation. You can manually change this port after installation.

**Important:** The client port number must be the same on both the management server and the standby management server in a high-availability relationship. If you change the client port number on either the management server or the standby management server, you must also change the port number on the other server.

1. Open the `rmserver.properties` file in the `TPC_install_directory\wlp\usr\servers\replicationServer\properties\` directory.
2. Modify the port number for the following property:  
`communications.port=port_number`
3. Open the `repli.properties` file in the `TPC_install_directory\cli\` directory.
4. Modify the port number for the following property:  
`port=port_number`
5. Restart Tivoli Storage Productivity Center for Replication. You must restart Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the Tivoli Storage Productivity Center for Replication management servers and must be maintained on each Tivoli Storage Productivity Center for Replication management server.

---

## Managing storage systems

To replicate data among storage systems using IBM Tivoli Storage Productivity Center for Replication, you must add connections to the storage systems. After a storage system is added, you can associate a location, modify connection properties, set volume protection, and refresh the storage configuration for that storage system.

### Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration. You can then omit storage systems for which Tivoli Storage Productivity Center for Replication does not manage replication and omit storage systems that are being managed by another Tivoli Storage Productivity Center for Replication management server.

For redundancy, you can connect a single storage system using a combination of direct, Hardware Management Console (HMC), and IBM z/OS connections.

You can use the following storage systems:

- IBM TotalStorage Enterprise Storage Server Model 800

- IBM System Storage DS6000
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize family
- IBM Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)
- IBM XIV Storage System

SAN Volume Controller and Storwize systems can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage other storage systems through a single SAN Volume Controller or Storwize cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize clusters.

You can define a location for each storage system and for each site in a session. When you are adding copy sets to the session, only the storage systems whose location matches the location of the site are allowed for selection. This ensures that a session relationship is not established in the wrong direction.

### **Easy Tier heat map transfer**

The IBM System Storage DS8000 Easy Tier heat map transfer function transfers heat map information from a source storage system to one or more target storage systems. Each target storage system then generates volume migration plans based on the heat map data and the physical configuration of the storage system.

This process ensures that the performance characteristics of the target storage systems are consistently updated to reflect the performance characteristics of the source storage system.

The Easy Tier heat map transfer function is available for System Storage DS8000 Release 7.1 and later.

You can enable heat map transfers for System Storage DS8000 storage systems from the IBM Tivoli Storage Productivity Center for Replication graphical user interface (GUI). The storage systems must meet the following requirements:

- The source and target storage systems must be connected to Tivoli Storage Productivity Center for Replication by using a Hardware Management Console (HMC) connection.
- The Easy Tier heat map transfer function must be enabled on the source and target storage systems. If the function is disabled on either the source or target storage system, an error is displayed when the next transfer is attempted.

To enable, disable, and configure heat map transfers, see “Configuring the use of the Easy Tier heat map transfer function” on page 330.

### **Important:**

When you enable or disable the use of the Easy Tier heat map transfer function in Tivoli Storage Productivity Center for Replication, the function is not enabled or disabled on the storage systems that are connected to Tivoli Storage Productivity Center for Replication. The configuration options that you set for Easy Tier heat map transfer in Tivoli Storage Productivity Center for Replication are used only by Tivoli Storage Productivity Center for Replication.

Do not manage heat map transfers for the same System Storage DS8000 storage systems from multiple Tivoli Storage Productivity Center for Replication servers. Doing so causes transfers to occur more frequently than the default transfer schedule, which is 12 hours after the previous transfer.

### Easy Tier heat map transfer log files

Easy Tier heat map transfer log files are in the *TPC\_install\_directory\wlp\usr\servers\replicationServer\logs\ethmt\log\* directory. If a problem related to the Easy Tier heat map transfer function occurs, you can send these log files to IBM Software Support for troubleshooting purposes.

## Storage connections

You must create a connection from the IBM Tivoli Storage Productivity Center for Replication management server to each storage system.

The following types of connection are available. The connection type that you can use depends on the storage system type.

### Direct Connection

This connection type is available for all storage systems other than IBM System Storage DS8700 or later storage systems or System Storage DS8000 storage systems that are on an Internet Protocol version 6 (IPv6) network.

### Hardware Management Console (HMC) Connection

This connection type is available for System Storage DS8700 or later storage systems and System Storage DS8000 storage systems that are on an IPv4 or IPv6 network.

### IBM z/OS Connection

This connection type is available for System Storage DS8000, System Storage DS6000, and IBM TotalStorage Enterprise Storage Server Model 800 storage systems. If the storage system is connected to a z/OS system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added. If the storage system is connected to another z/OS system, you can use an IP connection to connect to the system.

A single storage system can be connected by using multiple connections for redundancy. For example, you can connect to a System Storage DS8000 storage system by using an HMC connection and a z/OS connection.

When you add a storage connection to the Tivoli Storage Productivity Center for Replication configuration, the storage system and the connection are added to the active management server and the standby server configuration.

The storage systems are not required to be connected to the standby management server. However, if a storage system does not have a connection on the standby management server, you cannot manage copy services on the storage system from the standby server.

**Important:** If the Metro Mirror heartbeat is enabled, do not connect to a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage system by using both an HMC connection and a direct connection. If you have both types of connections and the direct connection is lost, the session changes to the suspended state even though the HMC

connection is still valid. If both connections are lost and the session is in the suspended state, restart the session when connectivity is regained to synchronize the session with the hardware.

### **Direct connection**

By using a TCP/IP connection, the Tivoli Storage Productivity Center for Replication management server can connect directly to certain storage systems. The TCP/IP connection is required to discover a system's configuration (such as LSSs, volumes, volume size, and format), issue queries, and receive asynchronous events.

The Tivoli Storage Productivity Center for Replication management server can connect directly with the following systems:

- TotalStorage Enterprise Storage Server
- DS6000
- DS8000
- SAN Volume Controller
- IBM Storwize family
- Storwize V7000 Unified
- The XIV system

DS8000 storage systems on an IPv4 network can be connected directly to the management server. A direct connection requires an Ethernet card in the cluster. DS8000 systems on an IPv6 network cannot use a direct connection. They can be connected only through an HMC or z/OS connection.

When you add a direct connection to a DS or ESS cluster, specify the following information for clusters 0 and 1:

- IP addresses or domain names
- Ports
- User names
- Passwords

SAN Volume Controller and Storwize systems can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage other storage systems through a single SAN Volume Controller or Storwize cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize clusters.

When you add a direct connection to a cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the cluster IP address of the cluster. Ensure that the user name and password are correct for the cluster. If incorrect values are used, significant communication problems can occur, such as never advancing to the Prepared state.

**Important:** The storage system user name must have privileges to maintain SSH keys. For information about troubleshooting Secure Shell connections to the storage system, see the Ethernet Connection Restrictions on SAN Volume Controller website at [www-01.ibm.com/support/docview.wss?uid=ssg1S1002896](http://www-01.ibm.com/support/docview.wss?uid=ssg1S1002896).

### **Hardware Management Console connection**

The IBM Tivoli Storage Productivity Center for Replication management server can connect to DS8000 storage systems through a Hardware Management Console (HMC). An HMC can have multiple DS8000 storage systems connected to it. When

you add an HMC to the IBM Tivoli Storage Productivity Center for Replication configuration, all DS8000 storage systems that are behind the HMC are also added. You cannot add or remove individual storage systems that are behind an HMC.

You can also add a dual-HMC configuration, in which you have two HMCs for redundancy. You must configure both HMCs identically, including the user ID and password.

If planned maintenance is necessary on the HMC, it is recommended that you disable the Metro Mirror heartbeat on the management server while the maintenance is performed.

If the HMC needs to go down frequently or reboots frequently, it is recommended that you disable the Metro Mirror heartbeat. If the Metro Mirror heartbeat is required, the direct connection is recommended instead of an HMC connection.

**Important:** If a DS8000 storage system uses an HMC connection, the Metro Mirror heartbeat could trigger a freeze on the storage system and impact applications for the duration of the long busy timeout timer if the HMC is shut down for any reason, including upgrading microcode. The long busy timeout timer is the time after which the storage system will allow I/O to begin again after a freeze occurs if no run command has been issued by Tivoli Storage Productivity Center for Replication. The default value is two minutes for CKD volumes or one minute for fixed block volumes.

**Notes:**

- The user ID that you use to connect to the HMC must have admin, op\_storage, or op\_copy\_services privileges on the DS8000 storage system. In addition, the user ID must have Copy Services access to any volumes that are in resource groups on the DS8000 storage system.
- Although you can use either an IP address or a host name to connect to an HMC, the use of a host name ensures that connection to the HMC is not lost if the IP address changes.
- You cannot modify the primary IP address or host name for the HMC in Tivoli Storage Productivity Center for Replication after you add the connection. If you want to change the IP address or host name, you must add a new connection to the HMC and then remove the old connection.
- For minimum microcode requirements to connect to a DS8000 through a management console, see the Supported Storage Products List website at [www-01.ibm.com/support/docview.wss?uid=swg21386446](http://www-01.ibm.com/support/docview.wss?uid=swg21386446).

**z/OS connection**

This connection type is available for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800 storage systems. If the storage system is connected to an IBM z/OS system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added. If the storage system is connected to another z/OS system, you must use an IP connection to connect to the z/OS system.

To add a z/OS connection, the system must be listed on the Host Connections page and in the Connected state. You can add a connection to a z/OS system from Tivoli Storage Productivity Center for Replication that is running on any operating system. For example, you can add a connection to a z/OS system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX. Or,

you can add a connection in Tivoli Storage Productivity Center for Replication that is running on z/OS to another z/OS system.

The z/OS connection is used to issue replication commands and queries for attached count key data (CKD) volumes over an existing Fibre Channel network and to receive asynchronous events. When a storage system is added to Tivoli Storage Productivity Center for Replication through a z/OS connection, all CKD volumes that are attached to the Tivoli Storage Productivity Center for Replication management system are added to the Tivoli Storage Productivity Center for Replication configuration. CKD volumes that are not attached to the Tivoli Storage Productivity Center for Replication z/OS management server are not added to the Tivoli Storage Productivity Center for Replication configuration through the z/OS connection.

Ensure that all volumes in the logical storage subsystem (LSS) that you want to manage through a z/OS connection are attached to a z/OS system. Either the entire LSS must be attached to z/OS or none of the volumes in the LSS should be attached to z/OS for Tivoli Storage Productivity Center for Replication to properly manage queries to the hardware. For example, if there are two CKD volumes in an LSS, and one volume is attached to Tivoli Storage Productivity Center for Replication by using a z/OS connection and the other is attached through a direct connection, Tivoli Storage Productivity Center for Replication issues commands to both volumes over the Fibre Channel network. However, commands issued to the direct connection volume fail, and Tivoli Storage Productivity Center for Replication shows that the copy set that contains that volume has an error.

**Tip:** It is recommended that you create both a Hardware Management Console (HMC) and direct connection for CKD volumes to allow for greater storage accessibility.

If a storage system was previously added to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection and later the storage system is added through a TCP/IP connection, all non-attached CKD volumes and fixed block volumes are added to the Tivoli Storage Productivity Center for Replication configuration.

### **The Metro Mirror heartbeat**

The Metro Mirror heartbeat is not supported through a z/OS connection. To use the Metro Mirror heartbeat, the storage systems must be added by using an HMC or direct connection. If the Metro Mirror heartbeat is enabled and a storage system is added through a direct connection and a z/OS connection, a suspend results if the direct connection becomes disconnected because there is no heartbeat through the z/OS connection.

### **Removing connections to a z/OS system**

If Tivoli Storage Productivity Center for Replication has multiple connections to a storage system, the order in which you remove the connections produces different results:

- If you remove direct and HMC connections first, the fixed-block and non-attached CKD volumes that are attached through these connection types are removed from the Tivoli Storage Productivity Center for Replication configuration.



- The remaining CKD volumes that are attached through the z/OS host connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS host connection is removed.
- If you remove the z/OS host connection first and there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.

## Protected volumes

You can mark volumes as protected if you do not want those volumes used for replication.

When a volume is marked as protected, you cannot include that volume in a copy set. This protection applies only to IBM Tivoli Storage Productivity Center for Replication.

You might want to protect a volume in the following instances:

- The volume contains data that you never want to be copied to another volume. For example, the volume is secure, but if the data is copied to an unsecured volume, the data could be read. For this reason, the volume should not be the source for a relationship.
- The volume contains data that you do not want to be overwritten. For this reason, the volume should not be the target of a relationship.

Only administrators can change the volume protection settings.

## Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

**Note:** To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A\_Location and a session with the location of A\_Location for its H1

role, changing the location of the storage system to a different location, such as B\_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A\_Location, the session's role location is not changed.

**Important:** Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

## Adding a storage connection

You must add connections to storage systems before you can use Tivoli Storage Productivity Center for Replication to manage data replication.

To add a storage system connection, you must belong to the Administrator role.

### Adding a connection to a System Storage DS8000 storage system by using an HMC connection

You must use a Hardware Management Console (HMC) connection to connect to a System Storage DS8000 storage system if the system is a DS8700 or later or is on an Internet Protocol version 6 (IPv6) network.

You can also optionally use an HMC connection if the storage system is on an IPv4 network.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS8000** icon.
4. On the Connection page of the Add Storage System wizard, click **HMC** and complete the following fields.

#### Primary HMC

Specify the following information about the primary HMC that is used to connect to the storage system.

#### IP Address/Domain Name

Enter the IP address or host name for the primary HMC. Although either an IP address or a host name are acceptable, the use of a host name ensures that connection to the HMC is not lost if the IP address changes.

You cannot modify the primary IP address or host name for the HMC in Tivoli Storage Productivity Center for Replication after you add the connection. If you want to change the IP address or host name, you must add a new connection to the HMC and then remove the old connection.

#### Username

Enter the user name for the HMC.

#### Password

Enter the password for the HMC.

#### Secondary HMC (Optional)

Specify information about the secondary HMC, if applicable. A secondary HMC is used for redundancy.

Both HMCs must be configured identically and must have the same user name and password.

**IP Address/Domain Name**

Enter the IP address or host name for the secondary HMC.

5. Click **Next** and follow the instructions in the wizard.

### **Adding a connection to a System Storage DS8000 storage system by using a direct connection**

You can connect to a System Storage DS8000 storage system by using a direct connection if the system is earlier than a DS8700 and is on an Internet Protocol version 4 (IPv4) network.

Direct connection requires that each of the storage system clusters contains an Ethernet card.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS8000** icon.
4. On the Connection page of the Add Storage System wizard, click **Direct Connect** and complete the following fields.

**Cluster 0**

Specify the following information about connection to cluster 0:

**IP Address/Domain Name**

Enter the IP address or host name for the cluster.

**Port** Enter the port number for the cluster. The default value is 2433.

**Username**

Enter the user name for the cluster.

**Password**

Enter the password for the cluster.

**Cluster 1**

Specify the following information about connection to cluster 1. The port, user name, and password for cluster 0 are automatically entered in the corresponding fields for cluster 1. You can edit this information, if required for your environment.

**IP Address/Domain Name**

Enter the IP address or host name for the cluster.

**Port** Enter the port number for the cluster. The default value is 2433.

**Username**

Enter the user name for the cluster.

**Password**

Enter the password for the cluster.

5. Click **Next** and follow the instructions in the wizard.

### **Adding a connection to a System Storage DS6000 or TotalStorage Enterprise Storage Server storage system**

You can connect to a System Storage DS6000 or TotalStorage Enterprise Storage Server storage system by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS6000** or **ES800** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for the storage system clusters:

#### **Cluster 0**

Specify the following information about connection to cluster 0:

##### **IP Address/Domain Name**

Enter the IP address or host name for the cluster.

**Port** Enter the port number for the cluster. The default value is 2433.

##### **Username**

Enter the user name for the cluster.

##### **Password**

Enter the password for the cluster.

#### **Cluster 1**

Specify the following information about connection to cluster 1. The port, user name, and password for cluster 0 are automatically entered in the corresponding fields for cluster 1. You can edit this information, if required for your environment.

##### **IP Address/Domain Name**

Enter the IP address or host name for the cluster.

**Port** Enter the port number for the cluster. The default value is 2433.

##### **Username**

Enter the user name for the cluster.

##### **Password**

Enter the password for the cluster.

5. Click **Next** and follow the instructions in the wizard.

## **Adding a connection to a SAN Volume Controller or Storwize storage system**

You can connect to a SAN Volume Controller or Storwize storage system by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **SAN Volume Controller; Storwize Family** for storage systems in the IBM Storwize family; or **Storwize V7000 Unified** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for the storage system cluster:

#### **Cluster IP / Domain Name**

Enter the IP address or host name of the storage system node.

#### **Username**

Enter the user name for the node.

#### **Password**

Enter the password for the node.

5. Click **Next** and follow the instructions in the wizard.

### Adding a connection to an XIV system

You can connect to a IBM XIV Storage System by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **XIV** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for a node on the storage system. When you specify the information for one node, the remaining nodes are discovered automatically.

#### IP Address/Domain Name

Enter the IP address or host name for a node in the storage system.

#### Username

Enter the user name for the node.

#### Password

Enter the password for the node.

5. Click **Next** and follow the instructions in the wizard.

### Adding a connection to a z/OS system

You can connect to storage systems that are connected to an IBM z/OS system regardless of the operating system on which Tivoli Storage Productivity Center for Replication is installed.

If a z/OS system is defined on the Host Connections page, you can add a connection to a z/OS from any operating system. For example, you can add a connection to a z/OS system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX.

Complete the following steps to add a z/OS system connection:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **z/OS Direct Connection** icon.
4. On the Connection page of the Add Storage System wizard, select the storage systems that you want to add.
5. Click **Next** and follow the instructions in the wizard.

## Removing a storage system

You can remove a storage system from the Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to remove a storage system.

Removing a storage system removes all volumes on that storage system from management server control. All copy sets with a volume on the removed storage system are removed from their respective sessions, making the target volume unrecoverable. All connections to the removed storage system are removed, and any storage systems sharing these connections are also removed.

Perform these steps to remove a storage system:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, select the storage system that you want to remove.

**Important:** All connections to this storage system will be removed, all volumes on the storage system will be removed from management server control, and all copy sets that have a volume on this storage system will be removed from their respective sessions, leaving the target volume unrecoverable. Any storage systems sharing these connections will be removed as well.

3. From the **Select Action** list, select **Remove Storage System**.
4. Click **Yes** to remove the storage system.

## Modifying the location of storage systems

You can associate a location with a storage system after a connection has been made to that storage system.

**Prerequisites:** You must have Administrator privileges to modify the location of a storage system.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A\_Location and a session with the location of A\_Location for its H1 role, changing the location of the storage system to a different location, such as B\_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A\_Location, the session's role location is not changed.

Perform these steps to modify the location of a storage system:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, change the location of the storage system by selecting a previously defined location from the drop-down list or type a new name in the table cell.

To disable site awareness, set the location to **None**.

**Note:** Locations are deleted from the drop-down list when there is no longer a storage system with an association to that location.

## Modifying storage connection properties

You can modify the connection properties for a storage system, including IP addresses, user name and password.

**Prerequisites:**

- You must have Administrator privileges to modify storage connection properties.
- The storage system must be in the Disconnected state to change most storage connection parameters. You can add a secondary HMC to an existing HMC connection without the HMC being disconnected.

A storage system can lose connection to the management server, for example, if a port is blocked by a firewall or the user name or password is changed on the storage system. If the storage system loses connection, you might need to modify

parameters (for example, user name or password) manually on the storage system, and then update the parameters in IBM Tivoli Storage Productivity Center for Replication.

Perform these steps to modify storage connection properties:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click the **Connections** tab.
3. Click the storage connection ID.
4. Modify the appropriate settings to match the settings for the storage system.
5. Click **Apply**.

## Refreshing the storage system configuration

You can refresh the storage system configuration to query the storage system for changes, such as which volumes are contained in an LSS. You might do this when you reconfigure a storage system and you want Tivoli Storage Productivity Center for Replication to be aware of the changes.

**Prerequisites:** You must have Administrator privileges to modify storage connection settings.

Perform these steps to refresh the storage configuration:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, select the storage system for which you want to refresh the configuration.
3. From the **Select Action** list, select **Refresh Configuration**.

## Setting volume protection

To ensure that data on a volume is not overwritten, you set its status to protected. Protected volumes are excluded from replication.

You must have Administrator privileges to change the protection setting of a volume.

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click **Volume Protection**. The Volume Protection wizard is opened.
3. The Volume Protection wizard, select a storage system and complete the following steps:
4. Optional: Depending on the type of storage system:
  - a. Select **All IO Groups** or a specific I/O group.
  - b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
  - c. Select **All Pools** or a specific pool.
5. Optional: In the **Volume** field, select a single volume.
6. Optional: In the **Volume Mask** field, enter a sequence of characters and wildcards that match user-defined or system-defined volume IDs. To protect a specific volume, enter the volume ID such as `ESS:2105.FCA57:VOL:1000`. To use a pattern to retrieve one or more volume IDs, you can enter a partial volume ID and use the wildcard character (\*) to represent zero or more characters. For example, to retrieve all volume IDs that contain the characters `FCA57`, you enter `*FCA57*`.

7. Click **Next**.
8. Verify the search results, and click **Next**.
9. Click **Select All** to protect all the volumes. Alternatively, select a check box next to the volumes that you want to protect.
10. Click **Next**.
11. Click **Finish**.

## Restoring data from a journal volume

This topic provides information on restoring data from a journal (J) volume. It describes how to restore data from journal volume used as part of a ESS/DS6000/DS8000 Global Mirror session, or as part of a Metro Global Mirror session, if data was corrupted on a host volume after you issued a **Recover** command. Following these steps will enable you to return to a consistent copy of the data on the host volume.

Perform the following steps to move the data from the journal volume back to the host volume:

**Note:** Follow these instructions only if you have already issued a **Recover** command to the site containing the journal volume. After the **Recover** command is issued, the journal volume will hold a copy of the consistent data at the time the command was issued.

1. Outside of IBM Tivoli Storage Productivity Center for Replication, using the DS8000 GUI /CLI, issue withdraw initiate background copy (issue a **rmflash -cp** command) on pairs containing the journal volume (for example, H2J2). This copies the remaining uncopied tracks from the host to the journal. Then, ensure all the Out of Sync (OOS) tracks reach zero.
2. Create a separate FlashCopy session either with IBM Tivoli Storage Productivity Center for Replication, or with the DS8000 GUI /CLI (issue a **mkflash** command with background copy), with the following conditions:
  - The journal volume (Jx) is the source volume.
  - The host volume (Hx or Ix if using a session with Practice capabilities) is the target volume.
  - x is the site the **Recover** command was issued to.

## Configuring the use of the Easy Tier heat map transfer function

You can enable Easy Tier heat map transfers for System Storage DS8000 storage systems that are connected to IBM Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).

The heat map transfer configuration options that are available in Tivoli Storage Productivity Center for Replication include adding storage systems for which you want to transfer heat map data and enabling or disabling the transfer of Easy Tier heat map data for those storage systems.

1. In the menu bar, click **Storage > Storage Systems**.
2. Click the **Easy Tier Heat Map Transfer** tab.
3. The following configuration options are available:

**Status** Indicates whether the transfer of Easy Tier heat map data is enabled or disabled.



**Enable Transfer**

Click to enable the transfer of Easy Tier heat map data, and then click **Yes** in the confirmation window. Easy Tier heat map data is copied from the storage systems that are listed on this page to the target storage systems that are associated with those storage systems.

**Disable Transfer**

Click to disable the transfer of Easy Tier heat map data, and then click **Yes** in the confirmation window. The storage systems that are listed on this page remain displayed with a status of **Inactive**.

**Add Storage Systems**

Click to add System Storage DS8000 storage systems that you want to include in the Easy Tier heat map transfer configuration. If there are available systems, a list of storage systems is displayed. Select the storage systems from this list. Only those storage systems that meet the following requirements are displayed:

- The storage system must be version 7.1 or later.
- The storage system must be connected to IBM Tivoli Storage Productivity Center for Replication by using a Hardware Management Console (HMC) connection.
- The storage system is not already displayed on this page.

**Select Action**

Select a storage system that is listed on this page and then select one of the following actions:

**Remove Storage System**

Click to remove a storage system from this page. The Easy Tier heat map transfer process from the storage system to the target storage systems is stopped.

**View Transfer Details**

Click to view detailed information about the transfer from the source storage system to its target storage systems. This information includes whether the transfer succeeded or failed and the date and time that the transfer started and stopped. This action is not available if the source storage system does not have target storage systems.

**Storage system details**

The following information is shown for each storage system:

**Storage System**

Lists each source storage system and its target storage systems, if available.

**Connection Status**

One of the following states is displayed for each storage system connection:

**Connecting**

Indicates that connection to the storage system is in progress.

**Connected**

Indicates that connection to the storage system is established.

### **Disconnected**

Indicates that connection to the storage system was lost. The connection can be lost for several reasons, including:

- The storage system is not available.
- The user name and password that is used to connect to the storage system was changed.
- There are firewall issues that prevent connection to the storage system.

### **Inactive**

Indicates that the transfer of Easy Tier heat map data is disabled.

### **Related concepts:**

“Easy Tier heat map transfer” on page 318

The IBM System Storage DS8000 Easy Tier heat map transfer function transfers heat map information from a source storage system to one or more target storage systems. Each target storage system then generates volume migration plans based on the heat map data and the physical configuration of the storage system.

### **Transfer Results for Storage System window**

Use the Transfer Results for Storage System window to view details about the transfer of Easy Tier heat map data to target storage systems.

The following information is displayed in this window. This information includes the number of successful and failed transfers and the last start and stop time that a transfer occurred.

#### **Target Storage System**

Select the storage system for which you want to view the transfer details.

#### **Attempts**

Shows the number of attempts that were made to transfer the Easy Tier heat map data to the target storage system.

#### **Successful Transfers**

Shows the number transfers that were successful.

#### **Failed Transfers**

Shows the number of transfers that failed.

#### **Last Transfer Start Time**

Shows the date and time that the last transfer started. If a transfer is in progress, **In Progress** is shown.

#### **Last Transfer Stop Time**

Shows the date and time that the last transfer stopped. If a transfer is in progress, **Unknown** is shown.

#### **Last Transfer Result**

Provides a message that explains the results of the last transfer. If a transfer is in progress, a message specifying that the transfer is in progress is shown.

#### **Next Transfer Time**

Shows the date and time that the transfer is scheduled to occur again. If a transfer is in progress, **Unknown** is shown. Transfers are automatically scheduled for 12 hours after the previous transfer.

---

## Managing host systems

A host system is an IBM z/OS or AIX system that connects to storage systems to enable certain replication features for those systems.

A connection to a z/OS host system is required if you want to enable z/OS features such as HyperSwap and hardened freeze in Tivoli Storage Productivity Center for Replication sessions.

If Tivoli Storage Productivity Center for Replication is installed on the host system, the host system connection is automatically added to Tivoli Storage Productivity Center for Replication. This connection is referred to as the native z/OS connection.

If Tivoli Storage Productivity Center for Replication is not installed on the host system, you must add the connection to the host system by using an IP connection. To use an IP connection, you must set certain Resource Access Control Facility (RACF®) settings on the z/OS host as described in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

A connection to an AIX host system is required if you want to use the Open HyperSwap feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site. For the swap to occur, the storage systems that contain the volumes must be attached to an AIX host and the AIX host must be connected to Tivoli Storage Productivity Center for Replication. The Open HyperSwap feature is optional for the Metro Mirror with Failover/Failback session type.

**Restriction:** Open HyperSwap is not supported for AIX host servers that are in a clustered environment such as IBM PowerHA® SystemMirror® for AIX.

You can add both z/OS and AIX host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. For example, you can add a connection to a z/OS host system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX.

### Adding a host system connection

You can add a connection to one or more host systems to the Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to add a host system connection.

Follow these steps to add a host system connection:

1. In the menu bar, click **Storage > Host Connections**.
2. On the Host Connections page, click **Add Host Connection**.
3. In the Add Host Connection window, enter the following information:

**Connection Type**

Select the host system type: z/OS or AIX.

**Host name / IP address**

Enter the host name or IP address for the host system.

**Port** Enter the port for the host system.

For z/OS, the port number must be the same as the port number that is specified for the HyperSwap management address space IOSHMCTL SOCKPORT parameter on the z/OS system. For more information about the SOCKPORT parameter, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

For AIX, use the default port number 9930 unless the port was modified in Subsystem Device Driver Path Control Module (SDDPCM).

**User name**

Enter the user name. The **User name** field is displayed only for z/OS host systems. The user name and password must be the same name and password that are specified in the Resource Access Control Facility (RACF) settings on the host system. For more information about the RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

**Password**

Enter the password. The **Password** field is displayed only for z/OS host systems.

4. Click **Add Host**.

The host system is displayed on the Host Connections page. If connection to the host system is successful, the status **Connected** is displayed. If the connection to the host system failed, the status **Disconnected** is displayed.

## Modifying a host system connection

You can modify host system connections that are in the Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to modify a host system connection and the connection must be in the **Disconnected** state.

Follow these steps to modify a host system connection:

1. In the menu bar, click **Storage > Host Connections**.
2. On the Host Connections page, click the link for the host system in the **Host System** column. Complete the information in the following fields. The fields that are displayed depend on the host system and the type of connection to the host system.

**Note:** All fields are presented for informational purposes; however, some fields are information-only fields.

**Host name or IP address**

For AIX host systems or IBM z/OS host systems that are connected by using an IP connection, enter the host name or IP address for the host system.

If Tivoli Storage Productivity Center for Replication is installed on the z/OS host system, **ZOS\_NATIVE\_CONNECTION** is displayed.

**Port** Enter the port for the host system.

For z/OS, the port number must be the same as the port number that is specified for the HyperSwap management address space IOSHMCTL SOCKPORT parameter on the z/OS system. For more information about the SOCKPORT parameter, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

For AIX, use the default port number 9930 unless the port was modified in Subsystem Device Driver Path Control Module (SDDPCM).

**Type** Shows the type of host system connection.

**System name**

Shows the name of the z/OS host system. If there is no connection to the host system, **Unknown** is displayed in this field.

**Sysplex name**

Shows the name of the z/OS sysplex that the host system is in or specifies that the system is not in a sysplex. If there is no connection to the host system, **Unknown** is displayed.

**User name**

Enter the user name for the z/OS host system. The user name and password must be the same name and password that are specified in the Resource Access Control Facility (RACF) settings on the host system. For more information about the RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

**Password**

Enter the user password for the z/OS host system. If you change the user name, you must change the password.

**Sessions**

Shows the session that is associated with the AIX host system. The session is displayed only if the Open HyperSwap feature is enabled for the session.

The session name that is displayed might not be a session that is defined on the current Tivoli Storage Productivity Center for Replication server. For example, another Tivoli Storage Productivity Center for Replication server was used to connect to the host and a session name was defined on that server, but was not defined on the current server.

3. Click **Update Host**.

The host system is displayed on the Host Connections page. If connection to the host system is successful, the status **Connected** is displayed. If the connection to the host system failed, the status **Disconnected** is displayed.

## Removing a host system connection

You can remove host system connections from the Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to remove a host system.

Removing a host system connection disables the ability to use features that are supported by the host such as Open HyperSwap or HyperSwap. Sessions that use the host system to provide these features can no longer communicate with the host and the features are disabled for the entire session.

Follow these steps to remove a host system connection:

1. In the menu bar, click **Storage > Host Connections**.
2. On the Host Connections page, select the host system connection that you want to remove.

3. From the **Select Action** list, select **Remove Host Connection**.
4. Click **OK** to remove the host system connection.

The host system is removed from the Host Connections page.

## Removing a session from an AIX host system connection

You can remove a session that is associated with an IBM AIX host system from the Tivoli Storage Productivity Center for Replication configuration. When you remove a session, the host no longer recognizes that the session is managing the volumes that are attached to that host. This function is to be used primarily for cleanup purposes.

**Prerequisites:** You must have Administrator privileges to remove a session association.

When a session has Open HyperSwap enabled, the session communicates with the host system and the host system stores an association to that session on the Tivoli Storage Productivity Center for Replication server. If the Tivoli Storage Productivity Center for Replication server that made the association becomes inaccessible, it might be necessary to clean up and remove the session association from a different Tivoli Storage Productivity Center for Replication server.

If a host system has an associated session, the session name is displayed in the **View/Modify Host Connection** window for the host system. To open this window, click the link for the host system in the **Host System** column on the Host Connections page.

Follow these steps to remove a session association from a host system connection:

1. In the menu bar, click **Storage > Host Connections**.
2. On the Host Connections page, select the host system connection that contains the session that you want to remove.
3. From the **Select Action** list, select **Remove Session Association**.
4. Click **OK** to remove the session association.

The session name is removed from the **Sessions** field in the **View/Modify Host Connection** window.

---

## Managing logical paths

Logical paths define the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path. To configure logical paths for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000, use the ESS/DS Paths panel in Tivoli Storage Productivity Center.

To configure partnerships for the following storage systems, use the graphical user interface (GUI) or command-line interface (CLI) for the storage system:

- SAN Volume Controller
- Storwize
- The XIV system

## Viewing logical paths

You can view all logical paths that are defined on an TotalStorage Enterprise Storage Server, System Storage DS8000, or System Storage DS6000 storage system.

Complete one of these procedures to view logical paths:

- From the ESS/DS Paths page of Tivoli Storage Productivity Center for Replication:
  1. In menu bar, click **Paths**.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems page:
  1. In menu bar, click **Storage > Storage Systems**.
  2. Select the storage system for which you want to view logical paths.
  3. From the **Select Action** list, select **View Paths**. The paths page is displayed with a list of defined logical paths.

## Adding logical paths

You can add logical paths for IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 storage systems.

Ensure that you have defined the appropriate storage systems on the Storage Systems page.

Perform these steps to add logical paths:

1. In menu bar, click **Paths**.
2. Click **Manage Paths**. The Path Management wizard is displayed.
3. From the drop-down boxes in the Path Management wizard, select the source storage system, source logical subsystem, target storage system, and target logical subsystem. Then, click **Next**.
4. From the drop-down boxes in the Path Management wizard, select the source port and target port and click **Add**. You can add multiple paths between the logical subsystems, or just one at a time. When you have made your selections, click **Next**.
5. Confirm your selections and click **Next**.
6. Verify the remaining wizard panels and click **Next**.
7. Click **Finish**.

## Adding logical paths by using a CSV file

You can create a comma separated (CSV) file to define logical paths. The CSV file specifies storage systems pairings and associated port pairings that are used for replication. Tivoli Storage Productivity Center for Replication uses the port pairings that are defined in the CSV file to establish logical paths.

You can define the following types of port pairing for the source and target storage systems:

- Port pairings for all logical subsystems that are in a session.
- Port pairings for a logical subsystem to logical subsystem mapping.
- Port pairings for all logical subsystems. You can define a single set of port pairings for the logical subsystems or you can define different pairings for all even logical subsystems and all odd logical subsystems.

Complete these steps to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths by using a CSV file:

1. Create a CSV file named `portpairings.csv` in the `TPC_install_directory\wlp\usr\servers\replicationServer\properties\` directory. You can create the CSV file in a spreadsheet such as Microsoft Excel or in a text editor. The properties directory contains an example CSV file that is named `portpairings.csv.sample`, which contains the following example port pairs:

```
# 1. Each line in the file represents a storage system pairing
# and its associated port pairings.
# There are multiple formats that define
# different port pairings types:
#
# Specific Session: This format can be used if you
# want to use a different set of port pairings for each session.
# This may be necessary in some environments in order to prevent
# the paths of of session to affect another.
#
# Format:
# storageSystemA:storageSystemB:sessionName,
# port1:port2,...,portN:portM
#
# where storageSystemA and storageSystemB are the storage
# system IDs used in TPC-R to identify the storage systems
# and sessionName is the name of the session you wish this
# port pairing to be used for. port1, port2, portN, portM
# are hexadecimal numbers indicating the storage system
# communication ports.
#
# A valid example would be:
# 2107.04131:2107.01532:myPayRollSession,0x0331:0x0024,
# 0x0330:0x0100,0x0331:0x000C
#
# Specific LSS to LSS pairing: This format can be used
# if you want to specify that a specific LSS pairing use a
# specific set of port pairings.
#
# Format:
# storageSystemA:systemALSS:storageSystemB:systemBLSS,
# port1:port2,...,portN:portM
#
# where storageSystemA and storageSystemB are the storage system
# IDs used in TPC-R to identify the storage systems, and
# systemALSS/systemBLSS are LSSs on storage systems as in
# hexadecimal format. port1, port2, portN, portM
# are hexadecimal numbers indicating the storage system
# communication ports.
#
# A valid example would be:
# 2107.04131:0x52:2107.01532:0x5f,0x0331:0x0024,
# 0x0330:0x0100,0x0331:0x000C
#
# Box to Box pairing: This format can be used if you want
# to specify a set of port pairings that will be used for
# all LSS pairings across the same source and target box.
# This pairing has multiple options. A single set of port pairings
# can be defined for all LSS pairings between the boxes,
# OR the port pairings can define that the even source LSSs
# should use one set and the odd LSSs should use a different set.
# There are some scenarios particularly with Global Mirror
# support where performance can be improved by mapping
# even and odd LSSs.
#
# Format:
# Option 1:
# storageSystemA:storageSystemB,port1:port2,...,portN:portM
#
```



```

# where storageSystemA and storageSystemB are the storage system
# IDs used in TPC-R to identify the storage systems.
# port1:port2,... portN:portM are hexadecimal numbers indicating
# the storage system communication ports.
#
# A valid example would be:
# 2107.04131:2107.01532,0x0331:0x0024,0x0330:0x0100,0x0331:0x000C
#
# Option 2:
# storageSystemA:storageSystemB,even, port1:port2,...,portN:portM,
# odd,port1:port2,...portX:portY
#
# where storageSystemA and storageSystemB are the storage system IDs
# used in TPC-R to identify the storage systems.
# port1:port2,...portN:portM is the set of port pairings to use when
# the source is on an even LSS
# port1:port2,...portX:portY is the set of port pairings to use when
# the source is on an odd LSS
# All ports are hexadecimal numbers indicating the storage system
# communication ports.
#
# A valid example would be:
# 2107.04131:2107.01532,even,0x0331:0x0024,0x0330:0x0100,odd,0x0331:0x000C,
# 0x0334:0x0330
#

```

2. To enable the changes in the file, complete a task that requires new paths to be established. For example, suspend a session to remove the logical paths and then issue the **Start H1->H2** command to enable the paths to use the port pairings in the CSV file.

#### Considerations when creating and using the CSV file:

- The CSV file does not affect Global Mirror control paths.
- Port mapping is bidirectional. A logical path is established from system A to system B and from system B to system A depending on the direction of the pairs on the hardware.
- If the CSV file contains multiple lines that specify the same storage system to storage system pairing, Tivoli Storage Productivity Center for Replication uses the last line. This rule applies regardless of the order of the storage system pairing. For example, if you have storage systems 2107.04131:2107.01532 defined on the first line of the CSV file and then have 2107.01532:2107.04131 defined on the second line, Tivoli Storage Productivity Center for Replication uses second line.
- If a line in the CSV contains information that is not formatted correctly, the line is ignored. This rule includes lines that specify storage systems but do not include ports or include ports that are not formatted correctly.
- If the CSV file contains valid and invalid port pairs, the valid port pairs might or might not be established. Invalid port pairs can cause the following errors to be displayed in the Tivoli Storage Productivity Center for Replication console and on the ESS/DS Paths page:
  - Return Code F52: This error is displayed if a port is invalid.
  - Return Code 0400: This error is displayed if a port is invalid and out of the range for the device.

Other storage system error codes might be displayed also, depending on the path topology, types of paths, and the incorrect port pairings that are specified in the CSV file.

- If the CSV file contains no valid port pairs, no logical paths are established and subsequent commands to the storage systems that require logical paths might fail. If there are existing logical paths for a storage system, those paths are used until they are removed.

## Removing logical paths

You can remove logical paths that are defined on an TotalStorage Enterprise Storage Server, System Storage DS8000, or System Storage DS6000 storage system.

1. In menu bar, click **Paths**.
2. Click the link for the storage system that contains the paths that you want to remove.
3. Select the paths that you want to remove.
4. From the **Select Action** list, select **Remove**.

---

## Setting up data replication

This topic describes the how to set up data replication in your environment, including creating sessions and adding copy sets to those sessions.

A *session* is a container of multiple copy sets managed by a replication manager. A *copy set* is a set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size. In a replication session, the number of volumes in a copy set and the role that each volume in the copy set plays are determined by the session type.

## Sessions

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

The type of data replication that is associated with the session determines the actions that you can perform against all copy sets in the session, the number of volumes that a copy set can contain, and the role that each volume plays.

**Attention:** Use only the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) or command-line interface CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through other interfaces such as the System Storage DS CLI. If you modify relationships through other interfaces, a loss of consistency can occur across the relationships that are managed by the session. The exceptions to this requirement are failover operations that are managed by external applications for certain session and storage system types as described in “Failover and failback operations” on page 347.

### Single-target and multi-target sessions

Sessions are referred to as single-target or multi-target sessions. With single-target sessions, the source volume site can have only one target site. Data replication occurs from the source to the target.

With multi-target sessions, the source volume site can have multiple target sites. Data replication can occur from the source to an individual target or to all targets simultaneously.

**Related tasks:**

“Completing session administration tasks” on page 427

Use session commands to complete session tasks such as starting, stopping, suspending, or terminating sessions.

**Copy sets**

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

Each volume in a copy set must be of the same size and volume type. For example, SAN Volume Controller volumes must be used with other SAN Volume Controller volumes. The number of volumes in the copy set and the role that each volume plays is determined by the session type that is associated with the session to which the copy set belongs.

The following table lists the estimated number of role pairs and volumes per copy set that are supported for each session type.

*Table 80. Number of role pairs and volumes per copy set for each session type*

Session	Role Pairs	Volumes
Basic HyperSwap	1	2
FlashCopy	1	2
Snapshot <sup>1</sup>	0	1
Metro Mirror Single Direction	1	2
Metro Mirror Failover/Failback	1	2
Metro Mirror Failover/Failback with Practice	3	3
Metro Mirror - Metro Mirror	3	3
Global Mirror Single Direction (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	3	3
Global Mirror Single Direction (SAN Volume Controller, Storwize)	1	2
Global Mirror Failover/Failback (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	3	3
Global Mirror Failover/Failback (SAN Volume Controller, Storwize)	1	1
Global Mirror Failover/Failback with Practice (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	5	4
Global Mirror Failover/Failback with Practice (SAN Volume Controller, Storwize)	3	3
Global Mirror Either Direction with Two-Site Practice	8	6
Global Mirror Failover/Failback with Change Volumes	1	4
Metro Global Mirror	6	4
Metro Global Mirror with Practice	8	5

Table 80. Number of role pairs and volumes per copy set for each session type (continued)

Session	Role Pairs	Volumes
1. An XIV Snapshot session requires that the user define only the H1 volumes. All target volumes are created on the same storage pool as the source volumes.		

Use the Add Copy Sets wizard to add copy sets to a session. You can select a storage system; a logical subsystem (LSS), I/O group, or pool; or a single volume for each role and then create one or more copy sets for the session.

You can use one of the following volume pairing options to automatically create multiple copy sets in the same session.

**Storage system matching (System Storage DS8000, System Storage DS6000, or TotalStorage Enterprise Storage Server Model 800 Metro Mirror sessions only)**

Creates copy sets by matching volumes (based on the volume IDs) across all LSSs for the selected storage systems. For example, volume 01 on the source LSS is matched with volume 01 on the target LSS.

You can select the storage system and select **All Logical Subsystems** in the list of LSSs. You can also do auto-matching at the LSS level for Metro Mirror sessions.

**LSS, I/O group, or pool matching**

Creates copy sets by matching all volumes based on the selected LSS, I/O group, or pool for each role in copy set.

Select the storage system and LSS, I/O group, or pool, and then select **All Volumes** in the **Volume** list.

If you do not want to use the auto-generated volume pairing for a copy set, clear that copy set so that it is not added during the wizard. Then, add the remaining copy sets and reopen the Add Copy Set wizard and manually enter the volume pairings that you want.

Invalid copy set are not added to the session. Copy sets can be invalid if their volumes are not the same type or size.

You can remove copy sets that you do not want to add to the session, even if they are valid. This process enables you to filter and eliminate unwanted copy sets before they are added to the session.

You can export the copy sets to take a snapshot of your session at a particular point in time for backup purposes.

**Note:** You can copy an entire storage system only for Metro Mirror sessions.

**Adding copy sets**

When you create a copy set for a session, a warning is displayed if one or more of the selected volumes exist in another session. If the volume you selected is in another session, confirm whether you want to add the volume to this session. Ensure that you intend to have the volume in multiple sessions because having the volume in multiple sessions can create conflicts. A scenario in which you would want the same volumes in multiple sessions is one in which you create a session for normal replication and then create another session for a disaster recovery practice. You must use the same target volumes in both sessions.

You can use extent space-efficient volumes as copy set volumes depending on the capabilities of the storage system.

## Removing copy sets

You remove a copy set or range of copy sets by selecting the following items:

- Source volume
- LSS, I/O group, or pool
- Storage system

When the list of copy sets that meet your criteria is displayed, you can select the copy sets that you want to remove.

The consequence of removing copy sets varies depending on the state of the session:

### Defined

There is no relationship on the hardware. The copy set is removed from the Tivoli Storage Productivity Center for Replication data store.

### Preparing or Prepared

The copy set is currently copying data, so Tivoli Storage Productivity Center for Replication terminates the hardware relationship for the copy set. The rest of the copy sets continue to run uninterrupted.

### Suspended or Target Available

Any existing relationships on the hardware are removed for the copy set.

Before removing all copy sets from that session, terminate the session. Removing the copy sets when the session is active can considerably increase the amount of time it takes for the copy set removal to complete. Copy sets are removed one at a time, and when the session is active, commands are issued to the hardware. However, if you terminate the session first, then commands are not issued to the hardware and the removal process completes faster.

**Tip:** When you remove a copy set from Tivoli Storage Productivity Center for Replication, you might want to keep hardware relationships on the storage systems. These relationships are useful when you want to migrate from one session type to another or when resolving problems. For more information about keeping the hardware relationships when removing copy sets, see [Removing Copy Sets](#).

Depending on the storage system, the following actions occur when a copy set is removed:

### ESS 800, DS6000, and DS8000:

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- Any peer-to-peer remote copy (PPRC) pair that is part of a Global Mirror consistency group is removed from the consistency group on the storage system.
- If the PPRC pair is part of a Global Mirror consistency group and is the last remaining source volume in a subordinate session, the subordinate session is removed from the storage system.
- If the PPRC pair is the last remaining participant in a Global Mirror session, the Global Mirror session is removed from the storage system.
- Any PPRC relationship remains on the storage system.

- A Metro Mirror (synchronous PPRC) pair that is in a HyperSwap configuration is removed from that configuration but the pair remains on the hardware.
- FlashCopy relationship remains on the storage system if the hardware has not completed any background copy.

#### **SAN Volume Controller, Storwize, Storwize V7000 Unified, or the XIV system**

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- FlashCopy, Metro Mirror, and Global Mirror relationships are pulled out of their consistency group. If they are the last remaining relationship in a consistency group, that consistency group is removed from the hardware.

When you specify the force removal option, all knowledge of the specified copy set is removed from Tivoli Storage Productivity Center for Replication, even if the relationship itself still exists. In this situation, you are not able to remove the relationship by using Tivoli Storage Productivity Center for Replication, because no information about the relationship exists. If you force a removal of a copy set and the removal fails, you must manually remove the relationship from the hardware. If you do not, you cannot to create new relationships.

One benefit of forcing a removal of the copy sets is that Tivoli Storage Productivity Center for Replication does not manage the consistency of copy sets that it has no knowledge of. This means that additional commands to the session do not affect the removed copy sets, even though they are still in a relationship on the hardware.

If you do not specify the force removal option and an error occurs that prevents the hardware relationships from being removed, the copy set is not removed from Tivoli Storage Productivity Center for Replication. The copy set remains as part of the session, and you can still perform actions on it.

#### **Related tasks:**

“Creating sessions and adding copy sets” on page 414

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

“Removing copy sets” on page 433

This topic describes how to remove copy sets.

#### **Volume roles**

Volume roles are given to every volume in the copy set. The role defines how the volume is used in the copy set and the site location of the volume. For example, the H1 role is made up of host-attached volumes that are located at the primary site.

The site determines the location of the volumes. The number of sites in a copy set is determined by the session type. Tivoli Storage Productivity Center for Replication supports up to three sites:

- Site 1** The location of the primary storage that contain the source data. Upon initial configuration, this site contains the host volumes with updates that are copied to the target volumes.

**Site 2** The location of the secondary storage that receives the copy updates from the primary storage.

**Site 3** The location of the tertiary storage that receives data updates from the secondary storage.

The volume roles that are needed in a copy set are determined by the type of replication that is associated with the session. Tivoli Storage Productivity Center for Replication supports these volume roles:

**Host volume**

A host volume is a volume that is connected to a server that reads and writes I/O. A host volume can be the source of updated tracks when the server connected to the host volume is actively issuing read and write input/output (I/O). A host volume can also be the target of the replication. When the host volume is the target, writes are inhibited.

Host volumes are abbreviated as Hx, where *x* identifies the site.

**Journal volume**

A journal volume stores data that has changed since the last consistent copy was created. This volume functions like a journal and holds the required data to reconstruct consistent data at the Global Mirror remote site. When a session must be recovered at the remote site, the journal volume is used to restore data to the last consistency point. A FlashCopy replication session can be created between the host or intermediate volume and the corresponding journal volume after a recover request is initiated to create another consistent version of the data.

Journal volumes are abbreviated as Jx, where *x* identifies the site.

**Intermediate volume**

An intermediate volume receives data from the primary host volume during a replication with practice session. During a practice, data on the intermediate volumes is flash copied to the practice host volumes.

Depending on the replication method being used, data on intermediate volumes might not be consistent.

Intermediate volumes are abbreviated as Ix, where *x* identifies the site.

**Target volume**

A target volume receives data from a host or intermediate volume. Depending on the replication type, that data might or might not be consistent. A target volume can also function as a source volume. For example, a common use of the target volume is as a source volume to allow practicing for a disaster (such as data mining at the recovery site while still maintaining disaster recovery capability at the production site).

**Change volume (only SAN Volume Controller or Storwize storage systems)**

A volume that contains point-in-time images that are copied from the host or target volume.

Change volumes are abbreviated as Cx, where *x* identifies the site.

**Related concepts:**

“Session types” on page 57

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

“Copy sets” on page 341

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

## Role pairs

A *role pair* is the association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is shown using an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

*Participating role pairs* are role pairs that are currently participating in the session's copy.

*Non-participating role pairs* are role pairs that are not actively participating in the session's copy.

Snapshot sessions do not use role pairs.

## Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

**Note:** You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

## Consistency groups

For Global Mirror and Metro Global Mirror sessions, IBM Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that have been updated to preserve write order and are therefore recoverable.

*Data exposure* is the period when data is written to the storage at the primary site until data is replicated to storage at the secondary site. Data exposure is influenced by factors such as:

- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- Input/output (I/O) load on the storage systems

To manage data exposure, you can change the consistency group interval time. The *consistency group time interval* specifies how often a Global Mirror and Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A smaller interval



value causes the session to attempt to create consistency groups more frequently, which might also increase the processing load and message-traffic load on the storage systems.

## Failover and failback operations

If a session has *failover* and *failback* capabilities, you can perform a site switch in which you move the operation of Tivoli Storage Productivity Center for Replication from one site to another, and change the direction of the copy without completing a full copy.

Failover is the process of temporarily switching production to a backup facility following a scheduled maintenance period or a disaster at a production site. A failover operation can be followed by a failback operation, which is the process of returning production to its original location and resuming the copy between the two sites.

The failover and failback operations use remote mirror and copy functions for these tasks:

- To reduce the time that is required to synchronize volumes after switching sites during planned or unplanned outages.
- To eliminate the need to perform a full volume copy from the recovery site to the production site, which can reduce the time that is required to resume operations at the production site.

## Failover operations that are managed by other applications

Applications such as the IBM Series i Toolkit, VMware Site Recovery Manager, and Veritas Cluster Server can also manage failover operations for the session types that are listed in the following table.

Session type	Storage systems
Metro Mirror Single Direction	<ul style="list-style-type: none"> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>
Metro Mirror Failover/Failback	<ul style="list-style-type: none"> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> <li>• The XIV system</li> </ul>
Metro Mirror Failover/Failback with Practice	<ul style="list-style-type: none"> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>

Session type	Storage systems
Metro Mirror - Metro Mirror	<ul style="list-style-type: none"> <li>System Storage DS8000 (This session type is available only for System Storage DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.)</li> </ul>
Global Mirror Single Direction	<ul style="list-style-type: none"> <li>System Storage DS8000</li> <li>System Storage DS6000</li> <li>SAN Volume Controller</li> <li>Storwize V3700</li> <li>Storwize V7000</li> <li>Storwize V7000 Unified</li> </ul>
Global Mirror Either Direction with Two-Site Practice	<ul style="list-style-type: none"> <li>System Storage DS8000</li> <li>System Storage DS6000</li> </ul>
Global Mirror Failover/Failback	<ul style="list-style-type: none"> <li>System Storage DS8000</li> <li>System Storage DS6000</li> <li>SAN Volume Controller</li> <li>Storwize V3700</li> <li>Storwize V7000</li> <li>Storwize V7000 Unified</li> <li>The XIV system</li> </ul>
Global Mirror Failover/Failback with Practice	<ul style="list-style-type: none"> <li>System Storage DS8000</li> <li>System Storage DS6000</li> <li>SAN Volume Controller</li> <li>Storwize V3700</li> <li>Storwize V7000</li> <li>Storwize V7000 Unified</li> </ul>
Global Mirror Failover/Failback with Change Volumes	<ul style="list-style-type: none"> <li>SAN Volume Controller</li> <li>Storwize V3700</li> <li>Storwize V7000</li> <li>Storwize V7000 Unified</li> </ul>

If an application completes a failover operation for a session, the Severe status is displayed for the session on the Session Details page in Tivoli Storage Productivity Center for Replication and an error message is generated for the role pairs for which the failover occurred. To view the error message:

1. On the Session Details page, click the link for each role pair for which the failover occurred.
2. On the Role Pair Details page, click the **Show** link for each role pair. This link is in the **Details** column.

To change the session status to Normal, use the application to confirm that the role pairs that are in the session are consistent and then restart the session.

## Session types

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

The session types that are available depend on the storage system type and the edition of Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication that you are using as shown in this topic.

### Basic HyperSwap

Basic HyperSwap is a copy services solution for IBM z/OS version 1.9 and later. Basic HyperSwap provides high availability of data if a storage system failure occurs.

The Basic HyperSwap session type enables Tivoli Storage Productivity Center for Replication to manage Basic HyperSwap by using Metro Mirror or Metro Global Mirror replication.

Basic HyperSwap sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

*Table 81. Basic HyperSwap*

Supported editions	Supported storage systems
<ul style="list-style-type: none"><li>• Tivoli Storage Productivity Center for Replication Basic Edition for System z</li><li>• Tivoli Storage Productivity Center for Replication for System z</li></ul>	<ul style="list-style-type: none"><li>• TotalStorage Enterprise Storage Server Model 800</li><li>• System Storage DS8000</li><li>• System Storage DS6000</li></ul>

### FlashCopy

FlashCopy replication creates a point-in-time copy of the data on a source volume to a target volume. Data that existed on the target volume is replaced by the copied data.

FlashCopy sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 82. FlashCopy

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication Basic Edition for System z</li> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3500</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Snapshot

Snapshot replication creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Snapshot sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 83. Snapshot

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication Basic Edition for System z</li> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	The XIV system

## Metro Mirror Single Direction

Metro Mirror is synchronous replication that operates between two sites that are up to 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Metro Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are propagated to the target volume before the write operation finishes processing.

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site.

Metro Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 84. Metro Mirror Single Direction

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Metro Mirror Failover/Failback

Metro Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Metro Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 85. Metro Mirror Failover/Failback

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> <li>• The XIV system</li> </ul>

## Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice replication combines Metro Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 86. Metro Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Global Mirror Single Direction

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Global Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume.

With the Global Mirror Single Direction session type, Global Mirror replication is available only from the primary site.

Global Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 87. Global Mirror Single Direction

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice replication enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Either Direction with Two-Site Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

*Table 88. Global Mirror Either Direction with Two-Site Practice*

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> </ul>

## Global Mirror Failover/Failback

Global Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Global Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

*Table 89. Global Mirror Failover/Failback*

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> <li>• The XIV system</li> </ul>

## Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Global Mirror Failover/Failback with Change Volumes sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 90. Global Mirror Failover/Failback with Change Volumes

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• SAN Volume Controller Version 6.3 and later</li> <li>• Storwize V3700 Version 6.4.1 and later</li> <li>• Storwize V7000 Version 6.3 and later</li> <li>• Storwize V7000 Unified Version 1.3 and later (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice replication combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 91. Global Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)</li> </ul>

## Metro Global Mirror

Metro Global Mirror replication provides continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Metro Global Mirror replication enables you to switch the direction of the data flow so that you can use your secondary or tertiary site as your production site.

Metro Global Mirror sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.



Table 92. Metro Global Mirror

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800 (only H1 site)</li> <li>• System Storage DS8000</li> </ul>

## Metro Global Mirror with Practice

Metro Global Mirror with Practice replication combines Metro Mirror, Global Mirror, and FlashCopy capabilities to provide a point-in-time copy of the data on the tertiary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Global Mirror with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 93. Metro Global Mirror with Practice

Supported editions	Supported storage systems
<ul style="list-style-type: none"> <li>• Tivoli Storage Productivity Center for Replication for System z</li> <li>• All editions of Tivoli Storage Productivity Center</li> </ul>	<ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800 (only H1 site)</li> <li>• System Storage DS8000</li> </ul>

### Related concepts:

“Basic HyperSwap” on page 356

Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

“FlashCopy” on page 356

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

“Snapshot” on page 357

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

“Metro Mirror” on page 358

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is in one storage system and the target is in another storage system.

“Global Mirror” on page 365

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

“Metro Global Mirror” on page 370

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy

into a single session, where the Metro Mirror target is the Global Mirror source.

## Basic HyperSwap

Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare your system for disaster recovery and ensure high availability of data. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is frozen to ensure that it is consistent.

**Restriction:** This replication method is available only for ESS, DS6000, and DS8000 storage systems.

Basic HyperSwap replication performs the following actions:

- Manages CKD volumes in Metro Mirror to manage synchronous peer-to-peer remote copy (PPRC) relationships.
- Permits only CKD volumes to be added to the HyperSwap session. The graphical user interface (GUI) shows only CKD volumes when you add a copy set. The command-line interface (CLI) does not add a copy set if a fixed block volume is specified.
- Monitors events that indicate a storage device failed.
- Determines whether the failing storage device is part of a Metro Mirror synchronous PPRC pair.
- Determines the action to be taken from policy.
- Ensures that data remains consistent.
- Swaps the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can occur from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

## Example

Jane is using multiple DS8000 storage systems. The host applications run on a z/OS operating system and the z/OS environment is connected to the DS8000 storage systems. She has a site in Manhattan and a secondary DS8000 in Hoboken. While it is not required that she has a disaster recovery solution in effect, she does need to have a high-availability solution to keep her applications running all the time. Jane is worried that if a volume fails on the DS8000 in Manhattan, her database application might fail. Even a small downtime can be costly to Jane.

Jane uses a Basic HyperSwap session to mirror the data on the DS8000 in Manhattan to the secondary DS8000 in Hoboken. If a volume at the Manhattan site fails, Basic HyperSwap automatically directs application I/O to the mirrored volumes in Hoboken.

## FlashCopy

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

The ESS, DS6000, and DS8000 provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

All platforms can use local replication in which the source volume is located in one LSS or I/O group and the target volume is located in the same or another LSS or I/O group. Using the FlashCopy feature, you can reference and update the source volume and target volume independently.

The following figure illustrates a FlashCopy session.



### Example

Jane works for a bank. Jane uses a FlashCopy session to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers start batch processing. Jane uses the FlashCopy session to create checkpoint restarts for the batch processing if the batch processing fails. In the batch processing, the first step is to balance all international accounts and to make a FlashCopy point-in-time copy of the resulting data. The second step in the batch processing is to complete the international disbursements.

If the second step in the batch process fails, Jane can use the data from the FlashCopy session that was taken of the first step to repeat the second step. In this way, Jane does not have to begin the entire process again. Jane also uses a CLI script that performs a FlashCopy operation every night at 11:59 PM, and another script that quiesces the database. She backs up the data to tape on the target storage system, and then sends the tape to the bank's data facility for storage.

### Snapshot

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV system uses a snapshot session to create a large number of volume copies without affecting system performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV system snapshots, no storage capacity is used by the snapshot until the source volume or the snapshot is changed.

The following figure illustrates a Snapshot session.



### Example

Jane's host applications are using an XIV system for their back-end storage. With the XIV system, Jane can create a large number of point-in-time copies of the data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any point in time.

Jane sets up a Snapshot session by using Tivoli Storage Productivity Center for Replication and specifies the volumes on the XIV system that are used by the host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly configure a single Snapshot session.

When the session is configured, Jane uses a CLI script that runs a **Create Snapshot** command to the session every two hours. If a problem occurs, such as data becoming corrupted, Jane can find a snapshot of the data from a time before the problem occurred. She can restore the data to that point.

By creating a set of snapshots of the data, Jane can also schedule batch processing against that data every day. She can use the batch processing to analyze certain trends in the market without affecting the host applications.

### Metro Mirror

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is in one storage system and the target is in another storage system.

**Attention:** If you have sessions that contain Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless you select the **Manage H1-H2 with HyperSwap** or the **Enable Hardened Freeze** option for the session. By using these options, z/OS Input/Output Supervisor (IOS) manages freeze operations for the volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself.

Metro Mirror replication maintains identical data on both the source and target. When a write operation is issued to the source copy, the changes that are made to the source data are propagated to the target before the write operation finishes processing. If the storage system ends, no data is lost when you use Metro Mirror if data must be used from the recovery site.


A Metro Mirror session in Global Copy mode creates an asynchronous relationship to accommodate the high volume of data that is migrated. As a result, the data on

the target system might no longer be consistent with the source system. The Metro Mirror session switches back to a synchronous relationship when Metro Mirror reissues a **Start** command. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate periods in which you require host I/O response time over data consistency.

**Tip:** To determine whether there is any out-of-sync data that must to be copied before the session is consistent, check the percent that is complete in the session details page.

### **Metro Mirror Single Direction**


The following figure illustrates a Metro Mirror Single Direction session.



### **Metro Mirror Failover/Failback**

Using Metro Mirror Failover/Failback, the data exists on the second site, which is less than 300 KM away, and you can switch the direction of the data flow. You can use this session type to run your business from the secondary site, and to copy changes that are made at the second site back to the primary site when you want to resume production at the primary site.

The following figure illustrates a Metro Mirror with Failover/Failback session.




### **Metro Mirror Failover/Failback with Practice**

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy features to provide a point-in-time copy of the data on the remote site. You can use this session type to practice what you might do if a disaster occurs, without losing your disaster recovery capability.

This solution consists of two host volumes and an intermediate volume.

The following figure illustrates a Metro Mirror Failover/Failback with Practice session.




### Metro Mirror Failover/Failback with HyperSwap

You can enable a Metro Mirror Failover/Failback session to have HyperSwap capabilities. To enable HyperSwap processing, see “Setting up the environment for HyperSwap” on page 374.

Metro Mirror Failover/Failback with HyperSwap combines the high availability of Basic HyperSwap with the redundancy of a two-site Metro Mirror Failover/Failback solution for managing count key data (CKD) volumes on z/OS. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with minimal effect on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Session Actions** list in the graphical user interface (GUI) or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates a Metro Mirror Failover/Failback session that is enabled for HyperSwap.




### Metro Mirror Failover/Failback with Open HyperSwap

You can enable a Metro Mirror Failover/Failback session to have Open HyperSwap capabilities. To enable Open HyperSwap processing, see “Setting up the environment for Open HyperSwap” on page 375.

Metro Mirror Failover/Failback with Open HyperSwap combines the high availability of Basic HyperSwap on z/OS for fixed-block AIX volumes with the redundancy of a two-site Metro Mirror Failover/Failback solution. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with minimal effect on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Session Actions** list in the GUI or when you issue a `cmdsess -action hyperswap` command.

The following figure illustrates a Metro Mirror Failover/Failback session that is enabled for Open HyperSwap.




### **Metro Mirror - Metro Mirror**

A Metro Mirror - Metro Mirror session is a multi-target session that consists of three sites. You can define any of the sites as the primary site and then run Metro Mirror replication from the primary site to either of the other sites individually or both sites simultaneously. For example, if Site 1 is the primary site, data replication can occur between the H1 and H2 volumes and the H1 and H3 volumes separately or at the same time.

This session type is available only for System Storage DS8000 storage systems with a microcode level that supports single source to multi-target relationships. To determine whether you can use this session type, refer to the System Storage DS8000 documentation for the microcode level that you are using.

The following figure illustrates a Metro Mirror - Metro Mirror session.



## Examples

Read the following scenarios for information on using Metro Mirror for synchronous, remote data replication between two sites.

### Metro Mirror Single Direction

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for the critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

In a synchronous copy method, when a write is issued to change the source, the change is propagated to the target before the write is posted. This method of replication maintains identical data in both the source and target. The advantage of this method is when a disaster occurs, there is no data loss at the recovery site because both writes must complete before signaling completion of a write to the source application. Because the data must be copied to both System Storage DS8000 devices before the write is completed, Jane can be sure that the data is safe.

The night before the planned outage, Jane quiesces the database and servers in San Francisco and starts the database and servers in Oakland. To accomplish this task, Jane issues the **Suspend** and **Recover** commands, and then issues the **Start** command on the secondary site. She shuts down the equipment in San Francisco to avoid any power spikes when she restarts the system after the power is turned on.

### Metro Mirror in Global Copy mode

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage,



so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for the critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

Jane wants to limit the effect on any applications while completing the initial Metro Mirror synchronization, so she begins the session in Global Copy mode. After she sees that approximately 70% of the data is copied, Jane decides to switch the session to Metro Mirror mode, assuring data consistency.

### **Metro Mirror with Practice**

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify the recovery procedure for the Oakland site, but she cannot stop running the Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane can practice the disaster recovery scenario in Oakland while the Metro Mirror session runs uninterrupted. By practicing running the applications at the Oakland site, Jane is better prepared to recover data if a disaster occurs at the San Francisco site.

While her session is running in Prepared state, Jane practices a recovery at the Oakland site by issuing the **Flash** command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while the session continues to replicate data from San Francisco to Oakland. As a result, she can carry out the recovery testing without stopping the replication for any extended duration of time.

If, at some point, the Metro Mirror session is suspended because of a failure, Jane can use the practice session to restart the data replication process. She maintains a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a **Recover** command to create a consistent version of the data on the H2 volumes. After the **Recover** command completes, she can issue the **Start H1->H2** command to resynchronize the data from the San Francisco site to the Oakland site.

If a failure occurs before the restarted session is in Prepared state, she has a consistent version of the data on the H2 volumes. She only has to issue the **Recover** command to put the session into Target Available state and make the H2 volumes accessible from the servers. If the session was not in Prepared state when it was suspended, the subsequent **Recover** command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes is not overwritten if the data to be copied to them is not consistent.

### **Metro Mirror Failover/Failback enabled for Open HyperSwap**

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify the recovery procedure for the Oakland site, but she cannot stop running the Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane can practice the disaster recovery scenario in Oakland while the Metro Mirror session runs uninterrupted. By practicing running the applications at the Oakland site, Jane is better prepared to recover data if a disaster occurs at the San Francisco site.

While the session is running in a Prepared state, Jane practices a recovery at the Oakland site by issuing the **Flash** command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while the session continues to replicate data from San Francisco to Oakland. As a result, she can carry out the recovery testing without stopping the replication for any extended duration of time.

If the Metro Mirror session is suspended because of a failure, Jane can use the practice session to restart the data replication process while she maintains a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a **Recover** command to create a consistent version of the data on the H2 volumes. After the **Recover** command completes, she can issue the **Start H1->H2** command to resynchronize the data from the San Francisco site to the Oakland site.

If a failure occurs before the restarted session is in Prepared state, she has a consistent version of the data on the H2 volumes. She only has to issue the **Recover** command to put the session into Target Available state and make the H2 volumes accessible from the servers. If the session was not in Prepared state when it was suspended, the subsequent **Recover** command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes is not overwritten if the data to be copied to the volumes is not consistent.

### Selecting a HyperSwap session

A global insurance company decided to use Tivoli Storage Productivity Center for Replication to manage its disaster recovery environment. Jane wants minimal data exposure, both for planned outages such as routine maintenance, and for unplanned disasters. They have CKD volumes on System Storage DS8000 devices, and use z/OS operating systems. They have two data centers in New York.

Jane chooses a Metro Mirror recovery solution because her priority is to protect the system from regional disasters. Jane decides to use Metro Mirror solution because her company has two data centers that are located near each other. Jane realizes that because she uses a z/OS operating system, CKD, and System Storage DS8000 hardware, she can also use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize the effects on any applications, while she maintains the failover process to the secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is the best solution.

After installing and configuring Tivoli Storage Productivity Center for Replication on z/OS, Jane starts the Tivoli Storage Productivity Center for Replication GUI. She adds the Tivoli Storage Productivity Center for Replication storage devices that she intends to use on all sites. From the Sessions page, Jane opens the Create Session window and selects the **Metro Mirror Failover/Failback** session type. After completing the information in the window, Jane clicks **Launch Add Copy Sets Wizard**. After she completes the wizard, she selects the **Manage H1-H2 with HyperSwap** option in the View/Modify Properties notebook.

Jane then issues a **Start H1->H2** command. After the initial copy is completed, Jane can safely replicate the data between both sites. She can

also issue the **HyperSwap** command between sites 1 and 2 to switch sites with minimal effect on the application during either a disaster or maintenance period.

### **Performing a planned HyperSwap**

Jane's company used Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane must perform maintenance on an H1 volume. During this time, Jane does not want the applications or replication to be interrupted. To prevent this interruption, before the maintenance is scheduled to begin, Jane uses the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap operation to the H2 volumes. This process changes the applications so that the data is written to H2. To perform a planned HyperSwap operation, Jane issues a **HyperSwap** command.

### **Understanding what happens when an unplanned HyperSwap occurs**

Several weeks after the planned maintenance at Jane's company is completed, an incident occurs at the H1 site. A disk controller fails, causing one of the H1 volumes to encounter a permanent I/O error. Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and the H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application changes to write data to the H2 volumes. The applications are not interrupted.

Jane configured a Simple Network Management Protocol (SNMP) listener to alert her to any events, so she receives the SNMP event that indicates that a HyperSwap occurred. Jane investigates the cause of the HyperSwap process and uses the z/OS console to identify the volume that triggered the HyperSwap process. Jane replaces the faulty disk controller. Then, to recover from the unplanned HyperSwap process, Jane issues the **Start H2->H1** command.

## **Global Mirror**

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume. However, changes can occur on the source volume before the target volume verifies that it received the change.

Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If the system fails, Global Mirror might lose some data that was transmitted when the failure occurred.


### **Global Mirror Single Direction**

The Global Mirror Single Direction session type enables you to run Global Mirror replication from the primary site.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the single direction session consists of two host volumes and a journal volume. The following figure illustrates a Global Mirror Single Direction session for these storage systems.



For SAN Volume Controller and Storwize storage systems, each copy set in the Global Mirror Single Direction session consists of two host volumes. The following figure illustrates a Global Mirror Single Direction session for these storage systems.




### Global Mirror Failover/Failback

The Global Mirror Failover/Failback with Practice session type combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the Global Mirror Failover/Failback session consists of two host volumes and a journal volume. The following figure illustrates a Global Mirror Failover/Failback session for these storage systems.



For SAN Volume Controller and Storwize storage systems, each copy set in the Global Mirror Failover/Failback session consists of two host volumes. The following figure illustrates a Global Mirror Failover/Failback session for these storage systems.



### **Global Mirror Failover/Failback with Change Volumes**


Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Change volumes are composed of a source change volume and a target change volume that contain a point-in-time image of the data from the source and target volumes, respectively.

A FlashCopy operation occurs between the source volume and the source change volume. The frequency of the FlashCopy operation is determined by the cycle period. For information about the cycle period and how it is set, see “Global Mirror Failover/Failback with Change Volumes session properties” on page 472. The data on the source change volume is then replicated to the target volume, and finally to the target change volume.

Because the data that is replicated between sites contains point-in-time changes rather than all changes, a lower bandwidth link is required between the sites. However, when you use change volumes, your data exposure can increase. Therefore, you might want to include or exclude change volumes depending on your network traffic or business requirements.


This session type is available only for SAN Volume Controller and Storwize storage systems. The following figure illustrates a Global Mirror Failover/Failback with Change Volumes session for these storage systems.



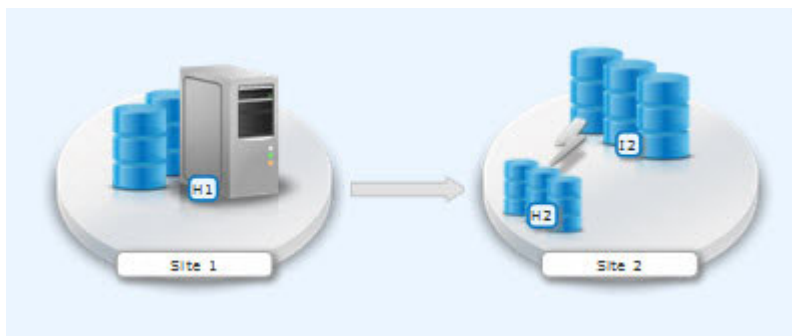
### Global Mirror Failover/Failback with Practice

The Global Mirror Failover/Failback with Practice session type combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the Global Mirror Failover/Failback with Practice session consists of two host volumes, an intermediate volume, and a journal volume. The following figure illustrates a Global Mirror Failover/Failback with Practice session for these storage systems.



For SAN Volume Controller and Storwize storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume. The following figure illustrates a Global Mirror Failover/Failback with Practice session for these storage systems.



## Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

This session type is available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems.

The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates a Global Mirror Either Direction with Two-Site Practice session for these storage systems.



### Examples

Read the following scenarios for information about using Global Mirror for asynchronous, remote data replication between two sites.

#### Global Mirror Single Direction

Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes a failure can occur, and she wants to ensure that the data is safe. Because Jane works in San Francisco, she wants her other site to be far away if a disaster occurs locally. Her other site is based in Houston. A minor earthquake occurs in San Francisco and power and communications both fails. Jane arranged for the data on customer accounts that recently opened or closed to be asynchronously copied in Houston, using Global Mirror. Jane risks losing the bytes of data that were being processed when the tremor disrupted the processing in San Francisco. However, she understands that it is a minor inconvenience when weighed next to the value of backing up her data to an area that is not prone to earthquakes.

#### Global Mirror with Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site, but she cannot stop running the Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane can practice her disaster recovery scenario in Houston while the Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site, Jane is better prepared to recover the data if a disaster ever strikes the San Francisco site.

#### Global Mirror Either Direction with Two-Site Practice

Jane wants to run a Global Mirror with Practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site,

but she cannot stop running the Global Mirror session while she takes time to practice a recovery. By using a Global Mirror Either Direction with Two-Site Practice session, Jane can practice her disaster recovery scenario in Houston while the Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site, Jane is better prepared to recover the data if a disaster ever strikes the San Francisco site.

Jane can use the Global Mirror Either Direction with Two-Site Practice session to run asynchronous consistent data replication from either the San Francisco site or the Houston site. She can practice the disaster-recovery tasks at the target site, regardless of the location of the current production site. Jane's business can run a consistent Global Mirror session from its Houston site back to San Francisco while it runs a production site at Houston.

## Setting up Global Mirror for Resource Groups on System Storage DS8000

If resource groups are defined on a System Storage DS8000, Global Mirror session IDs might be defined for some users. To determine which session IDs are valid, you must modify the `rmserver.properties` file and add the following property:

```
gm.master.sessionid.gm_role,session_name = xx
```

where `gm_role` is the role that has the master volume (for example, H1 in a Global Mirror failover/failback session), `session_name` is the name of the session that uses the session ID, and `xx` is the decimal number for the session ID.

**Tip:** System Storage DS8000 represents session IDs as a two-digit hexadecimal number. Use the decimal version of that number. For example, if you want a Global Mirror Failover/Failback session to use a session ID of 0F, the decimal number is 15, as shown in the following example:

```
gm.master.sessionid.H2.11194_wprac=15
```

### Metro Global Mirror

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

This replication method is available on only TotalStorage Enterprise Storage Server Model 800 and System Storage DS8000 storage systems. You can select ESS storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.

You can use both ESS and DS8000 volumes in the H1 volume role. If ESS and DS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync processing, and the ESS storage system performs a full copy operation. Because you cannot use the Incremental Resync function with ESS, a full copy is required when you change from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.



**Attention:** If you have Tivoli Storage Productivity Center for Replication for System z sessions that contain Metro Global Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless you select the **Manage H1-H2 with HyperSwap** or the **Enable Hardened Freeze** option for the session. By using these options, z/OS Input/Output Supervisor (IOS) manages freeze operations for the volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal effect on applications at the local site. This remote mirroring function works in combination with FlashCopy to provide a disaster-recovery solution that includes the following features:

- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages
- Data currency at the remote site with minimal time lag at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If Tivoli Storage Productivity Center for Replication runs on a z/OS operating system, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site by using the HyperSwap feature. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap operation, which transparently redirects application I/O to the auxiliary storage system. The Global Mirror relationship continues to run uninterrupted throughout this process. With this configuration, you can achieve almost a zero data loss at larger distances.


Using synchronous mirroring, you can switch from local site H1 to remote site H2 during a planned or unplanned outage. It also provides continuous disaster recovery protection of the H2 and H3 sites. You do not have to configure H2, if a switch from site H1 occurs. With this configuration, you can reestablish H2->H1->H3 recoverability while production continues to run at site H2. Additionally, this setup can reduce the workload on site H1.

**Important:** In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS storage system, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not available on the ESS storage system, full copies are performed in many of the transitions.

### **Metro Global Mirror**

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site.


The following figure illustrates a Metro Global Mirror session.



### **Metro Global Mirror with Practice**

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this session to practice what you might do if a disaster occurred without losing your disaster recovery capability.

The session consists of three host volumes, an intermediate volume, and a journal volume. The following figure illustrates a Metro Global Mirror with Practice session.



**Note:** In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS storage system, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not available on the ESS storage system, full copies are performed in many of the transitions.

## Examples

Read the following scenarios for information on using Metro Global Mirror for continuous, remote data replication between three sites.

### Metro Global Mirror

Although Jane works in San Francisco, she wants to run the business from either Oakland, the secondary site, or Houston, the tertiary site. Jane can use Metro Global Mirror with Failover/Failback to switch the direction of the data flow, so that she can run the business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss at the secondary site and minimal data loss at the tertiary site.

### Metro Global Mirror with Practice

Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site. However, she cannot stop running the Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane can practice her disaster recovery scenario in Houston while the Metro Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site and is prepared to run the applications at the Oakland site, if necessary, Jane is prepared to recover data if a disaster occurs at the San Francisco site.

Jane can use Metro Global Mirror with Practice to switch the direction of the data flow, so that she can run the business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss at the secondary site and minimal data loss at the tertiary site.

## Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data if a primary disk storage system failure occurs. When a failure occurs in writing I/O to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and minimal application effect.

### Sessions that can be enabled for HyperSwap or Open HyperSwap:

You can create sessions that enable swapping, which provides a session with a highly available business continuity solution.

### Sessions that can enable HyperSwap

The following session types can enable HyperSwap processing:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror
- Metro Global Mirror with Practice

To enable HyperSwap processing, the following conditions must apply:

- The volumes that are managed by the session are only TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 storage systems volumes.
- The volumes are count key data (CKD) volumes that are attached to an IBM z/OS system.
- The z/OS system is listed on the Tivoli Storage Productivity Center for Replication Host Connections page and is in the Connected state.

### Sessions that can enable Open HyperSwap

Only the Metro Mirror with Failover/Failback session type can enable Open HyperSwap processing.

To enable Open HyperSwap processing, the following conditions must apply:

- The volumes in the session are System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
  - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
  - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server are established.

### Setting up the environment for HyperSwap:

You must set up an environment in which you can run HyperSwap processing before you can enable HyperSwap for a Tivoli Storage Productivity Center for Replication session.

To enable HyperSwap processing, the following conditions must exist:

- The HSIB and IOSHSAPI address spaces are started on the IBM z/OS system to which the storage systems are connected.

- All reserve volumes are converted to global enqueues (ENQs).
- The storage system volumes are count key data (CKD) volumes that are attached to a z/OS system. The volumes can be attached to the same z/OS system or different systems.
- If the storage system volumes are not connected to a z/OS system on which Tivoli Storage Productivity Center for Replication is installed, the required Resource Access Control Facility (RACF) settings are set on the z/OS system and the **SOCKPORT** parameter is defined for the HyperSwap management address space IOSHMCTL. This set up enables Tivoli Storage Productivity Center for Replication to connect to the z/OS system by using an IP connection. For the required RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Complete the following tasks to set up an environment in which you can run HyperSwap processing:

1. Add a connection for each z/OS system in the Tivoli Storage Productivity Center for Replication GUI. To add a connection, open the Host Connections page and click **Add Host Connection**. If the z/OS system is the system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added in the GUI. If the z/OS system does not have Tivoli Storage Productivity Center for Replication installed, you must use an IP connection to connect to the z/OS system.
2. Add a connection to the storage systems that are attached to the z/OS systems. To add a connection, open the Storage Systems page and click **Add Storage Connection**. You can connect to a storage system by using a Hardware Management Console (HMC), direct, or z/OS connection.
3. Create Basic HyperSwap, Metro Mirror with Failover/Failback, Metro Global Mirror, or Metro Global Mirror with Practice sessions and enable HyperSwap. To create a session, open the Sessions page and click **Create Session**.

### Setting up the environment for Open HyperSwap:

You must set up an environment in which you can run Open HyperSwap processing before you can enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

Complete the following steps before you enable Open HyperSwap.

1. Ensure that the IBM AIX hosts and IBM System Storage DS8000 meet the following hardware and software requirements:

#### AIX requirements

You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>. Click the link for the applicable release under **Agents, Servers, and Browsers**.

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX)

#### System Storage DS8000 hardware requirements

The Open HyperSwap function requires System Storage DS8000 5.1 or later.

You cannot use Open HyperSwap with High Availability Cluster Multi-Processing (HACMP™).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts by following the instructions at “Adding a host system connection” on page 333.
3. Assign copy set volumes from the storage device to the host by using the System Storage DS8000 command-line interface (CLI) or the graphical user interface (GUI).
4. Run the AIX `cfgmgr` command to discover the volumes that are assigned to the host.

#### **Considerations for Open HyperSwap and the AIX host:**

- A single session that has the Open HyperSwap function enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the session.
- For AIX 5.3, a single host can manage a maximum of 1024 devices that are enabled for Open HyperSwap processing on the host. A maximum of eight logical paths can be configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that are enabled for Open HyperSwap processing on the host. A maximum of 16 logical paths can be configured for each copy set in the session.
- If an application on the host opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be ended. The **Terminate** command fails. To end the session, you must either close the application or remove the copy sets from the session. If you remove copy sets from the session, you must ensure that the application writes the data to the correct volume when the copy set relationship is restored.
- Open HyperSwap processing can fail on a subset of hosts for the session and continue to run on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing data to volumes on the source system and to volumes on the target system.
- To enable Open HyperSwap processing on the host, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

#### *Configuring timers to support Open HyperSwap:*

You can configure timeout values for the storage system, Tivoli Storage Productivity Center for Replication, and AIX hosts systems. These values can affect the operation of Open HyperSwap.

The following various timeout values can affect Open HyperSwap:

#### **Storage system quiesce timeout value**

This value is the time when the storage system starts a quiesce operation. When the timer value expires, I/O is resumed on the primary device. The default timeout value is 2 minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, use the `chdev` command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

#### **Storage system long busy timeout value**

This value is the time in seconds that the logical subsystem (LSS) consistency group volume stays in the long busy state after a remote mirror and copy error is reported.

### **Timeout values for the applications that are on the host**

The various applications that are running on the host have timeout values. The timeout values vary depending on the application.

### **Considerations for setting timers**

Consider the following information for setting timers:

- If the host quiesce timer is set to a shorter value than the Tivoli Storage Productivity Center for Replication response timer, an I/O swap failure can occur. If a storage system triggers an unplanned failover and if the storage system quiesce timer expires before Tivoli Storage Productivity Center for Replication responds, the host attempts to write I/O to the primary volume where the loss of access occurred. If the hardware condition that caused the loss of access continues, the write I/O operation fails again and an unplanned Open HyperSwap is not performed.
- If the host quiesce timer is set to a longer value than the Tivoli Storage Productivity Center for Replication response timer, an application timeout might occur if Open HyperSwap takes too long to complete.

### **Enabling HyperSwap or Open HyperSwap for a session:**

Enabling HyperSwap or Open HyperSwap for a session provides business recovery and business continuity.

Ensure that you can use the HyperSwap or Open HyperSwap functions in your environment. See “Setting up the environment for HyperSwap” on page 374 or “Setting up the environment for Open HyperSwap” on page 375.

To enable HyperSwap or Open HyperSwap for a session, complete the following steps:

1. In the Tivoli Storage Productivity Center for Replication menu bar, click **Sessions**.
2. On the Sessions page, select the session.
3. From the **Session Actions** list, select **View/Modify > Properties**.
4. In the **View/Modify Properties** notebook, select the following properties on the applicable role pair options tab:

#### **For HyperSwap**

##### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

##### **Manage H1-H2 with HyperSwap**

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

##### **On Configuration Error:**

##### **Partition the system(s) out of the sysplex**

Select this option to partition a new system out

of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

**Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

**On Planned HyperSwap Error:**

**Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

**Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

**On Unplanned HyperSwap Error:**

**Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

**Requirement:** If you select this option, you must restart the system .

**Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

**For Open HyperSwap**

**Manage H1-H2 with Open HyperSwap**

Select this option to trigger an Open HyperSwap operation for volumes that are attached to an IBM AIX host. This option redirects application I/O to the target volumes when there is a failure on the host accessible volumes.

Tivoli Storage Productivity Center for Replication uses Open HyperSwap to manage the H1-H2 sequence of a Metro Mirror session.

Only volumes that are attached to host systems that are defined in the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

**Restarting an AIX Host System that is enabled for Open HyperSwap:**

When an AIX host system is restarted, the host automatically opens any volumes for I/O that were open before the system was restarted. If Open HyperSwap was



enabled for a set of volumes on the host system, the host must determine which storage system is the primary system before the host can allow the volumes to be opened.

If the Metro Mirror relationship for the set of volumes is in Prepared or Suspended state and the host has connectivity to both the primary and secondary storage systems, the host can determine through the hardware which storage system is the primary system. In this situation, the host automatically opens the volumes.

If the Metro Mirror relationship for the set of volumes is in Prepared state and the host has connectivity to only the secondary storage system, all I/O to the volumes might be blocked on the host system. The I/O might be blocked until the host can verify the primary volume in the relationship. The AIX varyonvg command fails to open the volumes for I/O to prevent the application from writing to an incorrect volume. If the host can determine which volume is the primary volume in the relationship and connectivity to the primary storage system is still lost, a Hyperswap event is triggered. This event causes all I/O to be automatically opened and directed to the secondary storage system.

If the Metro Mirror relationship for the set of volumes is in Target Available state or the host system does not have connectivity and is unable to determine which site is the primary site, all I/O to the volumes are blocked on the host system. The Target Available state occurs after a Hyperswap or a Recover command is issued for the session. The varyonvg command fails to open the volumes for I/O to prevent the application from writing to an incorrect volume.

#### **Unblocking I/O on the host system after a host system restart**

When any of the previous scenarios cause I/O to be blocked, manual actions might be necessary to remove the block.

If the relationships are in Target Available state on the hardware, issue the **Start** command to the session in the required direction of the relationship. This action defines the primary storage system for the host. The host system can allow the volumes to be opened for I/O processing.

If the relationships cannot be restarted, or the host cannot determine the primary storage system, it might be necessary to manually decouple the volumes on the host system.

To decouple the volumes, complete one of the following tasks:


- Stop the session or remove the copy set. This action requires a full copy when the relationships are restarted.
- Remove Object Data Manager (ODM) entries by using the following command:  
`odmdelete -o Volume_Equivalency`

**Restriction:** This command deletes copy set information; use it only for this scenario.

#### **Planned and unplanned swaps:**

After a session is enabled for HyperSwap or Open HyperSwap processing and reaches Prepared state, Tivoli Storage Productivity Center for Replication loads the configuration of volumes that can be swapped to a z/OS or AIX operating system.

When the load is complete, the session can start a planned or unplanned swap. A hyperswap symbol is displayed over the arrow that shows replication between sites, as shown in the following figure.



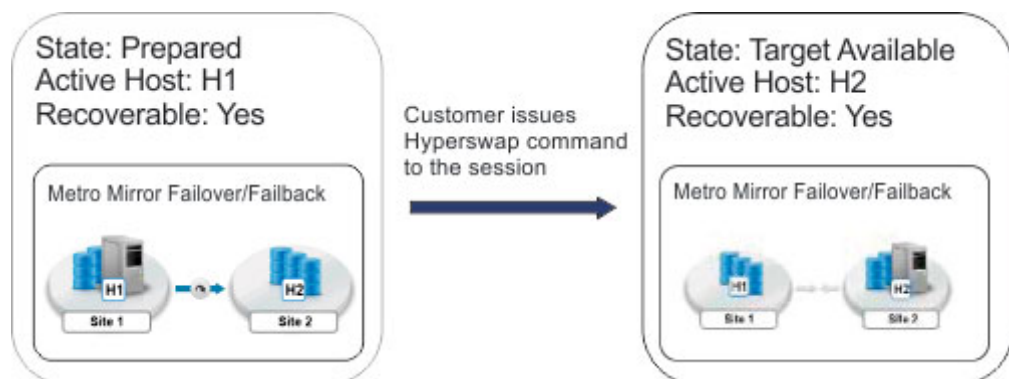
### Completing a Planned Swap

When the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. There may be cases such as a planned maintenance or a migration from the primary storage, in which a planned swap might be required. When the session is in swap capable state, a planned swap can be completed by issuing the **HyperSwap** command for the session.

After you run a planned swap for z/OS HyperSwap and Open HyperSwap, the session transitions to Target Available state and all the H1-H2 pairs are in Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal, which indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session has a status of SEVERE. The active host on the session is shown as H2.

All I/O is redirected to the H2 volumes. After a successful swap to site 2, you cannot enable the copy function to site 2. The only way to restart the copy processing is by issuing the **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, correct the error that caused the swap to occur, and manually restart the session to begin copying data to the other site.

The following figure illustrates a planned swap operation.



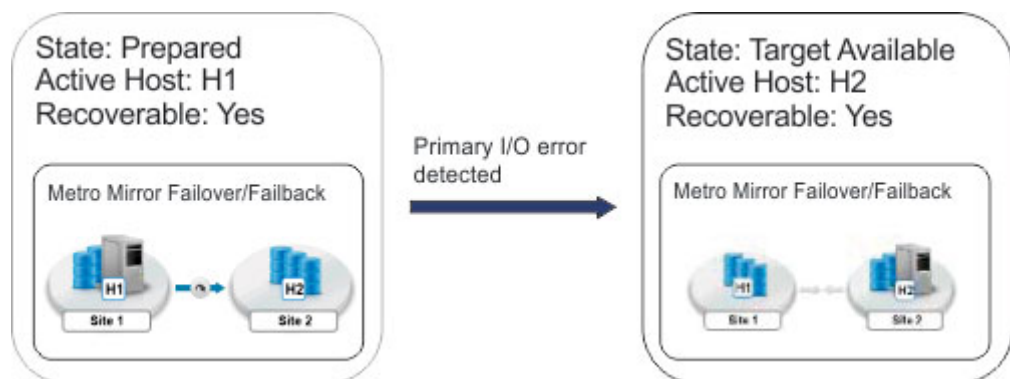
## Unplanned swaps

When the session configuration is loaded on z/OS for HyperSwap processing or on AIX for Open HyperSwap processing, the session is considered swap capable. If a primary I/O error occurs, a swap occurs automatically. For HyperSwap, z/OS completes the swap and then alerts Tivoli Storage Productivity Center for Replication that a swap occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to complete the swap.

When an unplanned swap occurs for HyperSwap and Open HyperSwap, the session transitions to Target Available state and all the H1-H2 pairs are in Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal and indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might show a status of SEVERE. The active host on the session is shown as H2.

All I/O is redirected to the H2 volumes. After a successful swap to site 2, you cannot enable the copy function to site 2. Therefore, it is not possible to issue a **Start H1->H2** command. The only way to restart the copy processing is by issuing a **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, correct the error that caused the swap to occur, and manually restart the session to begin copying data to the other site.

The following figure illustrates an unplanned swap.



*Verifying a session for a planned or unplanned swap:*

You can verify whether a session is capable of a planned or unplanned swap from the z/OS console on HyperSwap or the AIX host on Open HyperSwap.

Verify the status of HyperSwap from the z/OS console:

1. View the overall status of the HyperSwap session. Issue the `ds hs,status` command. For example:

```
15.03.06 SYSTEM1 d hs,status
15.03.06 SYSTEM1 STC00063 IOSHM0303I HyperSwap Status 531
Replication Session: SR_HS
HyperSwap enabled
New member configuration load failed: Disable
Planned swap recovery: Disable
Unplanned swap recovery: Disable
FreezeAll: No
Stop: No
```

- Verify all the volumes in the configuration. Issue the `ds hs,config(detail,all)` command. For example:

```
15.03.51 SYSTEM1 d hs,config(detail,all)
15.03.51 SYSTEM1 STC00063 IOSHM0304I HyperSwap Configuration 534
Replication Session: SR_HS
Prim. SSID UA DEV# VOLSER Sec. SSID UA DEV# Status
06 02 00F42 8K3602 06 04 00FA2
06 01 00F41 8K3601 06 03 00FA1
06 00 00F40 8K3600 06 02 00FA0
```

Verify the status of Open HyperSwap from the AIX host:

- View the session association and the path that the I/O is being routed to. The path is indicated by an asterisk. Issue the `pcmpath query device` command. For example:

```
host1> pcmpath query device 14

DEV#: 14 DEVICE NAME: hdisk14 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS Direction: H1<-H2
=====
PRIMARY SERIAL: 25252520000
-----
Path# Adapter/Path Name State Mode Select Errors
0 fscsi0/path0 CLOSE NORMAL 6091 0
1 fscsi0/path2 CLOSE NORMAL 6300 0
2 fscsi1/path4 CLOSE NORMAL 6294 0
3 fscsi1/path5 CLOSE NORMAL 6187 0

SECONDARY SERIAL: 34343430000 *
-----
Path# Adapter/Path Name State Mode Select Errors
4 fscsi0/path1 CLOSE NORMAL 59463 0
5 fscsi0/path3 CLOSE NORMAL 59250 0
6 fscsi1/path6 CLOSE NORMAL 59258 0
7 fscsi1/path7 CLOSE NORMAL 59364 0
```

### Temporarily disabling HyperSwap or Open HyperSwap:

In some situations, it might be necessary to temporarily disable the HyperSwap or Open HyperSwap function for a session.

You might want to disable HyperSwap or Open HyperSwap in the following circumstances:

- If you are performing maintenance on the system
- If one sysplex member cannot communicate with one or more volumes

Complete these steps to disable Open HyperSwap for a specific session:

- In the Tivoli Storage Productivity Center for Replication menu bar, click **Sessions**.
- On the Sessions page, select the sessions for which you want to disable the HyperSwap or Open HyperSwap function.
- From the **Session Actions** list, select **View/Modify > Properties**.
- On the applicable tab of the View/Modify Properties notebook, select **Disable HyperSwap** or **Disable Open HyperSwap**.

**Tip:** On management servers that run on a z/OS operating system, you can disable the HyperSwap function from an MVS command prompt window by entering the **SETHS DISABLE** command.

## Using active and standby Tivoli Storage Productivity Center for Replication servers with HyperSwap or Open HyperSwap:

You can set up an active and standby management server configuration with HyperSwap or Open HyperSwap.

### Active and standby servers with HyperSwap

When the storage system is set up to connect through the z/OS interface, the connection information is sent to the standby server and a connection is attempted. The connection might fail if the standby server does not run on a z/OS operating system or does not have access to the same volumes. If the connection fails and the standby server resumes the processing, the standby server cannot manage the HyperSwap processing.

On the z/OS system, if the session configuration was successfully loaded before the HyperSwap processing, the z/OS system can manage the HyperSwap processing. If the z/OS system swaps the volumes but cannot communicate with the Tivoli Storage Productivity Center for Replication server, the session is set to Suspended or Severe state. To set the session to Target Available state, clear the **Manage H1-H2 with Hyperswap** option, and issue the **Recover** command for the session.

### Active and standby servers with Open HyperSwap

When there is an active and standby management server configuration and a host system connection is added to the active server, a connection is attempted between the host system and the standby server. After AIX loads the session configuration, Open HyperSwap processing is possible only if there is communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a standby server resumes the processing and is unable to connect to the host system that manages the swap, the session cannot run Open HyperSwap processing. You must activate communication to the host system before the session can run Open HyperSwap processing.

#### Related tasks:

“Managing management servers” on page 308

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

## Session commands

The commands that are available for a session depend on the session type, and for some commands, the storage system type.

Commands are issued to Tivoli Storage Productivity Center for Replication sessions. Any subsequent command that is issued to an individual session is not processed until the first command is completed. Some commands, such as the **Start** command, can take an extended amount of time to complete because the commands set up the hardware relationships. You can continue to issue commands to other sessions during this time. When a command is completed, the console shows the results of the command.

### Basic HyperSwap commands

Use this information to learn about commands available for Basic HyperSwap sessions.

**Note:** Individual suspend and recover commands are not available in HyperSwap.

Table 94. Basic HyperSwap commands

Command	Action
<b>HyperSwap</b>	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Start H1-&gt;H2</b>	<p>Starts copying data synchronously from H1 to H2 in a Metro Mirror session.</p> <p><b>Note:</b> A session might go into a Severe state with error code 1000000 before the session returns to Normal/Prepared State and HyperSwap Capable. The duration of the Severe state depends on how large of a session is running.</p>
<b>Start H2-&gt;H1</b>	Starts copying data synchronously from H2 to H1 in a Metro Mirror session. You can issue this command only after the session has been swapped and the production site is H2. To enable data protection when the H1 volumes are available again, start I/O to the H2 volumes, and issue this command to replicate data from the H2 volumes to H1 volumes.
<b>Stop</b>	<p>Suspends updates to all the targets of pairs in a session. You can issue this command at any time during an active session.</p> <p><b>Note:</b> After you issue the stop command, targets might not be consistent.</p>
<b>Terminate</b>	Removes all physical copies and relationships from the hardware during an active session.

## FlashCopy commands

Use this information to learn about commands available for FlashCopy sessions.

Table 95. FlashCopy commands

Command	Action
<b>Flash</b>	<p>Performs the FlashCopy operation using the specified options. Issue the Flash command to create a data consistent point-in-time copy for a FlashCopy Session with volumes on the following storage systems:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V3500</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> <p>For a FlashCopy session containing ESS, DS6000, and DS8000 volumes, the Flash command by itself is not sufficient to create a consistent copy. To create a consistent copy using the ESS, DS6000, and DS8000 Flash commands, you must quiesce application I/O before issuing the Flash command.</p>
<b>InitiateBackgroundCopy</b>	<p>Copies all tracks from the source to the target immediately, instead of waiting until the source track is written to. This command is valid <i>only</i> when the background copy is not already running.</p>
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Start</b>	<p>Performs any steps necessary to define the relationship before performing a FlashCopy operation. For ESS, DS6000, and DS8000, this command is not an option. Issue this command to put the session in the prepared state for the following storage systems:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V3500</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>

Table 95. FlashCopy commands (continued)

Command	Action
<b>Terminate</b>	<p>Removes all active physical copies and relationships from the hardware during an active session.</p> <p>If you want the targets to be data consistent before removing their relationship, you must issue the <b>InitiateBackgroundCopy</b> command if NOCOPY was specified, and then wait for the background copy to complete by checking the copying status of the pairs.</p>

## Snapshot commands

Use this information to learn about commands that are available for Snapshot sessions and groups. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

Table 96. Snapshot session commands

Command	Action
<b>Create Snapshot</b>	Creates a snapshot of the volumes in the session
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Restore</b>	Restores the H1 volumes in the session from a set of snapshot volumes. You must have at least one snapshot group to restore from. When you issue this command in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI), you are prompted to select the snapshot group.

Table 97. Snapshot group commands

Command	Action
<b>Delete</b>	Deletes the snapshot group and all the individual snapshots that are in the group from the session and from the XIV system. If the deleted snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.



Table 97. Snapshot group commands (continued)

Command	Action
<b>Disband</b>	Disbands the snapshot group. When a snapshot group is disbanded, the snapshot group no longer exists. All snapshots in the snapshot group become individual snapshots that are no longer associated to the consistency group or the session. After a snapshot group is disbanded, it is no longer displayed in or managed by Tivoli Storage Productivity Center for Replication. If the disbanded snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.
<b>Duplicate</b>	Duplicates the snapshot group. When a snapshot group is duplicated, a new snapshot group is created with new snapshots for all volumes that are in the duplicated group. The name of the duplicated snapshot group is generated automatically by the XIV system.
<b>Lock</b>	Locks a snapshot group. If the snapshot group is locked, write operations to the snapshots that are in the snapshot group are prevented. By default, a snapshot group is locked when it is created. This action is valid only if the snapshot group is unlocked.
<b>Overwrite</b>	Overwrites the snapshot group to reflect the data that is on the H1 volume.
<b>Rename</b>	Renames the snapshot group to a name that you provide. The name can be a maximum of 64 alphanumeric characters.
<b>Restore</b>	Restores the contents of a snapshot group by using another snapshot group in the session. Both of the snapshot groups must contain the same subset of volumes.
<b>Set Priority</b>	Sets the priority in which a snapshot group is deleted. The value can be the number 1 - 4. A value of 1 specifies that the snapshot group is deleted last. A value of 4 specifies that the snapshot group is deleted first.
<b>Unlock</b>	Unlocks a snapshot group. If the snapshot group is unlocked, write operations to the snapshots that are in the snapshot group are enabled and the snapshot group is displayed as modified. This action is valid only if the snapshot group is locked.

## Metro Mirror commands

Use this information to learn about commands available for Metro Mirror sessions.

Table 98. Metro Mirror commands

Command	Action
<b>Enable Copy to Site 1</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available.
<b>Enable Copy to Site 2</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.

Table 98. Metro Mirror commands (continued)

Command	Action
<b>HyperSwap</b>	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Recover</b>	Issues the <b>Recover</b> command to suspended sessions. This command performs the steps necessary to make the target available as the new primary site. Upon completion of this command, the session becomes Target Available.
<b>Release I/O</b>	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the <b>Release I/O after Suspend</b> property is selected for the session.
<b>Start</b>	Establishes a single-direction session with the hardware and begins the synchronization process between the source and target volumes.
<b>Start H1-&gt;H2</b>	Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.
<b>Start H2-&gt;H1</b>	Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1. Indicates direction of a failover and failback between two hosts in a Metro Mirror session. If the session has been recovered such that the production site is now H2, you can issue the <b>Start H2-&gt;H1</b> command to start production on H2 and start data replication.
<b>Stop</b>	<p>Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session.</p> <p><b>Note:</b> Targets after the suspend are not considered to be consistent.</p>

Table 98. Metro Mirror commands (continued)

Command	Action
<b>StartGC</b>	Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2. While in the Preparing state, it will not change to the Prepared state unless you switch to Metro Mirror.
<b>Suspend</b>	<p>Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied.</p> <p><b>Note:</b> It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.</p> <p>When a <b>Suspend</b> command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.</p>
<b>Terminate</b>	Removes all copy relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the <b>Suspend</b> command, then the <b>Recover</b> command, and then the <b>Terminate</b> command.

## Metro Mirror with Practice commands

Use this information to learn about commands available for Metro Mirror with Practice sessions.

Table 99. Metro Mirror with Practice commands

Command	Action
<b>Enable Copy to Site 1</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available.
<b>Enable Copy to Site 2</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.
<b>Flash</b>	<p>Creates a FlashCopy image from I2 volumes to H2 volumes. The amount of time for this to occur will vary depending on the number of copy sets in the session.</p> <p><b>Note:</b> For ESS, DS6000, and DS8000 storage systems, the <b>Flash</b> command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.</p>

Table 99. Metro Mirror with Practice commands (continued)

Command	Action
<b>Recover</b>	<p>Takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2.</p> <p><b>Note:</b> The point-in-time copy is performed when the session is in a recoverable state to avoid that previous consistent data on H2 are overwritten by inconsistent data upon Recover. You can issue the <b>Flash</b> command if you want to force a point-in-time copy from I2 to H2 volumes afterward.</p>
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Release I/O</b>	<p>Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the <b>Release I/O after Suspend</b> property is selected for the session.</p>
<b>Start H1-&gt;H2</b>	<p>Establishes a Metro Mirror relationship from the H1 volumes to the I2 volumes, and begins data replication.</p>
<b>Start H2-&gt;H1</b>	<p>Establishes a Metro Mirror relationship from H2 to H1 and begins data replication.</p>
<b>StartGC_H1H2</b>	<p>Distinguishes when the session is in the Preparing state from H1 to I2 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.</p>
<b>StartGC_H2H1</b>	<p>Distinguishes when the session is in the Preparing state from H2 to H1 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.</p>

Table 99. Metro Mirror with Practice commands (continued)

Command	Action
<b>Suspend</b>	<p>Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied.</p> <p><b>Note:</b> The Metro Mirror command uses a freeze command on the ESS, DS6000, and DS8000 devices to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions will also become suspended. Avoid using the same LSS pairs for multiple Metro Mirror sessions.</p> <p>When a <b>Suspend</b> command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.</p>
<b>Stop</b>	<p>Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session.</p> <p><b>Note:</b> Targets after the suspend are not considered to be consistent.</p>
<b>Terminate</b>	Terminates all copy relationships on the hardware.
<b>Recover</b>	<p>Takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2.</p> <p><b>Note:</b> The point-in-time copy is performed when the session is in a recoverable state to avoid that previous consistent data on H2 are overwritten by inconsistent data upon Recover. You can issue the <b>Flash</b> command if you want to force a point-in-time copy from I2 to H2 volumes afterward.</p>

## Global Mirror commands

Use this information to learn about commands available for Global Mirror sessions.

Table 100. Global Mirror commands

Command	Action
<b>Enable Copy to Site 1</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available.
<b>Enable Copy to Site 2</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.

Table 100. Global Mirror commands (continued)

Command	Action
<b>Recover</b>	<p>Issue this command to recover the session to the target site. This command performs the steps necessary to make the target host volumes consistent and available for access as the new primary site. Upon completion of this command, the session state becomes Target Available. Do not access H2 volumes until the <b>Recover</b> command is completed and the session displays Target Available and Recoverable. A Recover to H2 also establishes a point-in-time copy to J2 to preserve the last consistent data.</p>
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Start</b>	<p>Establishes all relationships in a single-direction session and begins the process necessary to start forming consistency groups on the hardware.</p>
<b>Start H1-&gt;H2</b>	<p>Starts copying data from H1 to H2 in a Global Mirror Failover/Failback session. Establishes the necessary relationships in the session and begins the process necessary to start copying data from the H1 site to the H2 site and to start forming consistency groups.</p>
<b>Start H2-&gt;H1</b>	<p>Starts copying data from H2 to H1 in a failover and failback session for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000. If a recover was performed on a session such that the production site is now H2, you can issue a <b>Start H2-&gt;H1</b> command to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session. This action puts the relationships into a synchronized state and suspends them consistently. Sessions are consistent when a H2-&gt;H1 copy occurs for the following storage systems:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>

Table 100. Global Mirror commands (continued)

Command	Action
<p><b>StartGC H1-&gt;H2</b></p> <p>This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.</p>	<p>Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to H2. To change the session state from Preparing to Prepared, you must issue the <b>Start H1-&gt;H2</b> command and the session must begin to form consistency groups.</p> <p>There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of a primary storage system or a link failure between the sites occurs, the session might be inconsistent when you issue the <b>Recover</b> command.</p> <p>This command is available for the following storage systems:</p> <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS6000</li> <li>• System Storage DS8000</li> </ul>

Table 100. Global Mirror commands (continued)

Command	Action
<p><b>Suspend</b></p>	<p>Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H2 Global Copy pairs.</p> <p><b>If you are using a System Storage DS8000:</b> The <b>Suspend</b> command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level.</p> <p>If the following conditions are true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir -withsecondary</b>:</p> <ul style="list-style-type: none"> <li>• All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using.</li> <li>• The pause with secondary consistency option is enabled on the storage system.</li> </ul> <p>If these conditions are not true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir</b>.</p> <p>Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes.</p> <p>If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the <code>TPC_install_directory\wlp\usr\servers\replicationServer\properties\rmsserver.properties</code> file and add the following property:</p> <pre>csm.server.sus_gc_pairs_on_gm_pause = false</pre> <p><b>Tip:</b> The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the <code>csm.server.sus_gc_pairs_on_gm_pause</code> property to <b>false</b>.</p>
<p><b>Terminate</b></p>	<p>Removes all physical copies and relationships from the hardware during an active session.</p> <p>If you want the targets to be data consistent before you remove their relationship, you must issue the <b>Suspend</b> command, the <b>Recover</b> command, and then the <b>Terminate</b> command.</p>

## Global Mirror with Practice commands

Use this information to learn about commands available for Global Mirror with Practice sessions.



Table 101. Global Mirror with Practice commands

Command	Action
<b>Enable Copy to Site 1</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available.
<b>Enable Copy to Site 2</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.
<b>Flash</b>	<p>The Flash command ensures that all I2s are consistent, and then flashes the data from I2 to the H2 volumes. After the flash is complete, the Global Mirror session is automatically restarted, and the session begins forming consistency groups on I2. You can then use the H2 volumes to practice your disaster recovery procedures.</p> <p>For Global Mirror with Practice sessions for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000, the Flash command temporarily pauses the formation of consistency groups.</p>
<b>Recover</b>	Restores consistent data on I2 volumes and takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2. The I2 volumes continue to hold the consistent data and can be flashed again to H2 by using the <b>Flash</b> command.
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Start H1-&gt;H2</b>	Starts copying data from H1 to H2. After the first pass of the copy is complete for all pairs, the session establishes the I2->J2 FlashCopy pairs. The session starts the Global Mirror master so that the hardware begins to form consistency groups, to ensure that consistent data is at site 2.

Table 101. Global Mirror with Practice commands (continued)

Command	Action
<p><b>Start H2-&gt;H1</b></p>	<p>Starts copying data from H2 to H1 in a failover and failback session. If a recover was performed on a session such that the production site is now H2, you can issue a <b>Start H2-&gt;H1</b> command to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance.</p> <p><b>Note:</b> TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 volumes are not consistent for the <b>Start H2-&gt;H1</b> command. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend command to the session. The relationships are put into a synchronized state and are suspended consistently.</p>
<p><b>StartGC H1-&gt;H2</b></p> <p>This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.</p>	<p>Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to I2. To change the session state from Preparing to Prepared, you must issue the <b>Start H1-&gt;H2</b> command and the session must begin to form consistency groups.</p> <p>There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of the primary Tivoli Storage Productivity Center for Replication server occurs, the session might be inconsistent when you issue the <b>Recover</b> command.</p> <p>This command is available for the following storage systems:</p> <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS6000</li> <li>• System Storage DS8000</li> </ul>

Table 101. Global Mirror with Practice commands (continued)

Command	Action
<p><b>Suspend</b></p>	<p>Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H2 Global Copy pairs.</p> <p><b>If you are using a System Storage DS8000:</b> The <b>Suspend</b> command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level.</p> <p>If the following conditions are true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir -withsecondary</b>:</p> <ul style="list-style-type: none"> <li>• All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using.</li> <li>• The pause with secondary consistency option is enabled on the storage system.</li> </ul> <p>If these conditions are not true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir</b>.</p> <p>Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes.</p> <p>If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the <code>TPC_install_directory\wlp\usr\servers\replicationServer\properties\rmsserver.properties</code> file and add the following property:</p> <pre>csm.server.sus_gc_pairs_on_gm_pause = false</pre> <p><b>Tip:</b> The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the <b>csm.server.sus_gc_pairs_on_gm_pause</b> property to <b>false</b>.</p>
<p><b>Terminate</b></p>	<p>Removes all physical copies and relationships that are on the hardware.</p>

## Metro Global Mirror commands

Use this information to learn about commands available for Metro Global Mirror sessions.

Table 102. Metro Global Mirror commands

Command	Action
<b>Enable Copy to Site 1</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1-&gt;H3</b> command becomes available.
<b>Enable Copy to Site 2</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2-&gt;H3</b> command becomes available.
<b>HyperSwap</b>	Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
<b>RecoverH1</b>	Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H2->H1->H3) can move production to either the H1 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore, the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can set up for the failback.
<b>RecoverH2</b>	Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing is different depending on the recovery site. Therefore the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.
<b>RecoverH3</b>	Specifying H3 makes the H3 volume TargetAvailable. Metro Global Mirror (when H1->H2->H3 is run) can then move production to the H3 set of volumes. Because Tivoli Storage Productivity Center for Replication processing differs depending on the recovery site, the site designation is added to the <b>Recover</b> command so that Tivoli Storage Productivity Center for Replication can prepare for the failback.  This command prepares H3 so that you can start the application on H3. H3 becomes the active host, and you then have the option start H3->H1->H2 to perform a Global Copy copy back. The recovery establishes point-in-time copy to J3 volumes to preserve the last consistent data.
<b>Re-enable Copy to Site 1</b>	After you issue a <b>RecoverH1</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
<b>Re-enable Copy to Site 2</b>	After you issue a <b>RecoverH2</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
<b>Re-enable Copy to Site 3</b>	After you issue a <b>RecoverH3</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.

Table 102. Metro Global Mirror commands (continued)

Command	Action
<p><b>Refresh States</b></p>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<p><b>Release I/O</b></p>	<p>Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the <b>Release I/O after Suspend</b> property is selected for the session.</p>
<p><b>Start H1-&gt;H2-&gt;H3</b></p> <p>This command is the Metro Global Mirror initial start command.</p>	<p>Establishes Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this relationship includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.) The <b>Start H1-&gt;H2-&gt;H3</b> command can be used from some Metro Global Mirror configurations to transition back to the starting H1-&gt;H2-&gt;H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, target available, or suspended state.</p>
<p><b>Start H1-&gt;H3</b></p>	<p>From the H1-&gt;H2-&gt;H3 configuration, this command changes the session configuration to a Global Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. It can be used when a session is in preparing, prepared, and suspended states because there is not a source host change involved.</p> <p>You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and loses the near-zero data loss protection that is achieved with Metro Mirror when H1-&gt;H2-&gt;H3 is run. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a Global Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1-&gt;H3 configuration or from the TargetAvailable H2-&gt;H1-&gt;H3 state.</p>

Table 102. Metro Global Mirror commands (continued)

Command	Action
<p><b>Start H2-&gt;H3</b></p>	<p>From the H1-&gt;H2-&gt;H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This session is a host-volume change so this command is valid only when you restart the H1-&gt;H3 configuration or from the TargetAvailable H2-&gt;H1&gt;H3 state.</p> <p>From the H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, this command can be used when the session is in the preparing, prepared, and suspended states. The <b>Start H2-&gt;H1-&gt;H3</b> command can be used to transition back to the starting H2-&gt;H1-&gt;H3 configuration.</p>
<p><b>Start H2-&gt;H1-&gt;H3</b></p> <p>This command is the Metro Global Mirror start command.</p>	<p>This configuration completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, the J3 volume completes the Global Mirror configuration.</p>
<p><b>Start H3-&gt;H1-&gt;H2</b></p>	<p>After a recovery to H3 is completed, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long-distance copy back to site 1.</p> <p>To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a <b>Rcover</b> followed by a <b>Start H1-&gt;H2-&gt;H3</b> to go back to the original configuration.</p>

Table 102. Metro Global Mirror commands (continued)

Command	Action
<p><b>SuspendH2H3</b></p>	<p>Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H2 to H3 Global Copy pairs.</p> <p><b>If you are using a System Storage DS8000:</b> The <b>Suspend</b> command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level.</p> <p>If the following conditions are true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir -withsecondary</b>:</p> <ul style="list-style-type: none"> <li>• All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using.</li> <li>• The pause with secondary consistency option is enabled on the storage system.</li> </ul> <p>If these conditions are not true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir</b>.</p> <p>Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes.</p> <p>If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the <code>TPC_install_directory\wlp\usr\servers\replicationServer\properties\rmsserver.properties</code> file and add the following property:</p> <pre>csm.server.sus_gc_pairs_on_gm_pause = false</pre> <p><b>Tip:</b> The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the <b>csm.server.sus_gc_pairs_on_gm_pause</b> property to <b>false</b>.</p>
<p><b>SuspendH1H3</b></p>	<p>Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H3 Global Copy pairs.</p> <p>The information that is specific to System Storage DS8000 in the previous description for the <b>SuspendH2H3</b> command is also applicable to this command.</p>

## Metro Global Mirror with Practice commands

Use this information to learn about commands available for Metro Global Mirror with Practice sessions.

Table 103. Metro Global Mirror with Practice commands

Command	Action
<b>Enable Copy to Site 1</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session.
<b>Enable Copy to Site 2</b>	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session.
<b>Flash</b>	<p>This command is available in the following states:</p> <ul style="list-style-type: none"> <li>• Target Available state when the active host is H3  <b>Note:</b> Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a <b>Flash</b> command can be issued to complete the flash of the consistent data from I3 to H3.</li> <li>• Prepared state when the active host is H1 and data is copying H1 to H2 to I3, or the active host is H2 and data is copying H2 to H1 to H3.</li> <li>• Prepared state when the active host is H2 and data is copying H2 to I3.</li> <li>• Prepared state when the active host is H1 and data is copying H1 to I3.</li> </ul> <p>Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a <b>Flash</b> command can be issued to complete the flash of the consistent data from I3 to H3.</p> <p>Issuing a <b>Flash</b> command on a Global Mirror Practice session for ESS, DS6000, and DS8000 temporarily pauses the formation of consistency groups. Ensure that all I3s are consistent, and then flash the data from I3 to the H3 volumes. After the flash is complete, the Global Mirror session will be automatically restarted, and the session will begin forming consistency groups on I3. You can then use the H3 volumes to practice your disaster recovery procedures.</p>
<b>HyperSwap</b>	Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
<b>RecoverH1</b>	Specifying H1 makes the H1 volume TargetAvailable. When H1->H2->H3 is run, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback. The FlashCopy creates a consistent copy of the data on the H3 volumes so that an application can recover to those volumes and begin writing I/O. When the FlashCopy is complete, the session reaches a Target Available state, and you can attach your volumes on Site 3.



Table 103. Metro Global Mirror with Practice commands (continued)

Command	Action
<b>RecoverH2</b>	Specifying H2 makes the H2 volume TargetAvailable. When H1->H2->H3 is run, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.
<b>RecoverH3</b>	Specifying H3 makes the H3 volume the TargetAvailable. When H1->H2->H3 is run, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site; therefore, the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.
<b>Re-enable Copy to Site 1</b>	After you issue a <b>RecoverH1</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
<b>Re-enable Copy to Site 2</b>	After you issue a <b>RecoverH2</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
<b>Re-enable Copy to Site 3</b>	After you issue a <b>RecoverH3</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Release I/O</b>	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the <b>Release I/O after Suspend</b> property is selected for the session.
<b>Start H1-&gt;H2-&gt;H3</b>  This command is the Metro Global Mirror initial start command.	<p>This command creates Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this relationship includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.) The <b>Start H1-&gt;H2-&gt;H3</b> command can be used from some Metro Global Mirror configurations to return to the starting H1&gt;H2&gt;H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, target available, or suspended state.</p>

Table 103. Metro Global Mirror with Practice commands (continued)

Command	Action
<p><b>Start H1-&gt;H3</b></p>	<p>From the H1-&gt;H2-&gt;H3 configuration, this command changes the session configuration to a Global-Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source host change involved, it can be used when a session is in preparing, prepared, and suspended states.</p> <p>You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is completed, the session is running Global Mirror from H1 to H3 and loses the near-zero data loss protection that is achieved with Metro Mirror when H1-&gt;H2-&gt;H3 is run. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a Global-Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1-&gt;H3 configuration or from the TargetAvailable H2-&gt;H1-&gt;H3 state.</p>
<p><b>Start H2-&gt;H3</b></p>	<p>From the H1-&gt;H2-&gt;H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1&gt;H3 configuration or from the TargetAvailable H2-&gt;H1-&gt;H3 state.</p> <p>From the H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, this command can be used when the session is in the preparing, prepared, and suspended states. The <b>Start H2-&gt;H1-&gt;H3</b> command can be used to return to the starting H2-&gt;H1-&gt;H3 configuration.</p>
<p><b>Start H2-&gt;H1-&gt;H3</b></p> <p>This command is the Metro Global Mirror start command.</p>	<p>This configuration completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, the J3 volume completes the Global Mirror configuration.</p>

Table 103. Metro Global Mirror with Practice commands (continued)

Command	Action
<b>Start H3-&gt;H1-&gt;H2</b>	<p>After a recovery to H3 is completed, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long-distance copy back to site 1.</p> <p>To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a <b>Rcover</b> followed by a <b>Start H1-&gt;H2-&gt;H3</b> to go back to the original configuration.</p>
<b>SuspendH2H3</b>	<p>Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H2 to H3 Global Copy pairs.</p> <p><b>If you are using a System Storage DS8000:</b> The <b>Suspend</b> command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level.</p> <p>If the following conditions are true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir -withsecondary</b>:</p> <ul style="list-style-type: none"> <li>• All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using.</li> <li>• The pause with secondary consistency option is enabled on the storage system.</li> </ul> <p>If these conditions are not true, the <b>Suspend</b> command invokes a command that is equivalent to the System Storage DS8000 command <b>pausegmir</b>.</p> <p>Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes.</p> <p>If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the <code>TPC_install_directory\wlp\usr\servers\replicationServer\properties\rmsserver.properties</code> file and add the following property:</p> <pre>csm.server.sus_gc_pairs_on_gm_pause = false</pre> <p><b>Tip:</b> The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the <code>csm.server.sus_gc_pairs_on_gm_pause</code> property to <b>false</b>.</p>

Table 103. Metro Global Mirror with Practice commands (continued)

Command	Action
<b>SuspendH1H3</b>	<p>Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H3 Global Copy pairs.</p> <p>The information that is specific to System Storage DS8000 in the previous description for the <b>SuspendH2H3</b> command is also applicable to this command.</p>
<b>Terminate</b>	This command terminates all copy relationships on the hardware.

## Metro Mirror multi-target commands

Use this information to learn about commands available for Metro Mirror multi-target sessions.

Table 104. Metro Mirror multi-target command

Command	Action
<b>Confirm Production at Site 1</b>	Confirms that you want to establish site 1 as the production site after you issue suspend and recover commands for the session.
<b>Confirm Production at Site 2</b>	Confirms that you want to establish site 2 as the production site after you issue suspend and recover commands for the session.
<b>Confirm Production at Site 3</b>	Confirms that you want to establish site 3 as the production site after you issue suspend and recover commands for the session.
<b>FailoverH1</b>	Completes a failover operation to the H1 target site. This command enables read and write access to the target volumes, while maintaining a recording of changes to the volumes that prevents a full copy when the relationship is restarted. You can use this command to practice your disaster recovery capabilities without initiating a site switch. To switch sites, the RecoverH1 command is required.
<b>FailoverH2</b>	Completes a failover operation to the H2 target site. This command enables read and write access to the target volumes, while maintaining a recording of changes to the volumes that prevents a full copy when the relationship is restarted. You can use this command to practice your disaster recovery capabilities without initiating a site switch. To switch sites, the RecoverH2 command is required.
<b>FailoverH3</b>	Completes a failover operation to the H3 target site. This command enables read and write access to the target volumes, while maintaining a recording of changes to the volumes that prevents a full copy when the relationship is restarted. You can use this command to practice your disaster recovery capabilities without initiating a site switch. To switch sites, the RecoverH3 command is required.
<b>HyperSwap H1</b>	Triggers a HyperSwap operation where I/O is redirected from the source volume to the target volume H1, without affecting the application using those volumes.
<b>HyperSwap H2</b>	Triggers a HyperSwap operation where I/O is redirected from the source volume to the target volume H2, without affecting the application using those volumes.
<b>HyperSwap H3</b>	Triggers a HyperSwap operation where I/O is redirected from the source volume to the target volume H3, without affecting the application using those volumes.

Table 104. Metro Mirror multi-target command (continued)

Command	Action
<b>RecoverH1</b>	Recovers the session to the H1 target site. This command makes the target host volumes consistent and available for access as the new production site. Upon completion of this command, the session becomes Target Available.
<b>RecoverH2</b>	Recovers the session to the H2 target site. This command makes the target host volumes consistent and available for access as the new production site. Upon completion of this command, the session becomes Target Available.
<b>RecoverH3</b>	Recovers the session to the H3 target site. This command makes the target host volumes consistent and available for access as the new production site. Upon completion of this command, the session becomes Target Available.
<b>Re-enable Original Direction</b>	Sets the production site back to the original site and re-enables the start commands in the original direction. You can use this command after you issue the suspend, recover, and confirm commands.
<b>Refresh States</b>	<p>Refreshes the states of the role pairs that are in the session (if applicable for the session type) and refreshes the state of the session if it is incorrect. This command queries the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself.</p> <p>This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.</p> <p>This command is not available if the session is in the Defined state.</p>
<b>Release I/O</b>	Enables the source volume to receive data after a copy relationship is suspended. This command is not displayed if the <b>Release I/O after Suspend</b> property is selected for the session.
<b>Set Production to Site 1</b>	Establishes Site 1 as the production site while the session is in the Defined state. After you issue this command, the <b>Start H1-&gt;H2</b> and <b>Start H1-&gt;H3</b> commands become available.
<b>Set Production to Site 2</b>	Establishes Site 2 as the production site while the session is in the Defined state. After you issue this command, the <b>Start H2-&gt;H1</b> and <b>Start H2-&gt;H3</b> commands become available.
<b>Set Production to Site 3</b>	Establishes Site 3 as the production site while the session is in the Defined state. After you issue this command, the <b>Start H3-&gt;H1</b> and <b>Start H3-&gt;H2</b> commands become available.
<b>Start</b>	Establishes Metro Mirror relationships between the volumes on the production site and the volumes at the other sites in the session, and begins data replication to those sites. For example, if Site 1 is the production site, data replication occurs between the H1 and H2 volumes and the H1 and H3 volumes at the same time.
<b>Start H1-&gt;H2</b>	Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.

Table 104. Metro Mirror multi-target command (continued)

Command	Action
<b>Start H1-&gt;H3</b>	Establishes Metro Mirror relationships between the H1 volumes and the H3 volumes, and begins data replication from H1 to H3.
<b>Start H2-&gt;H1</b>	Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1.
<b>Start H2-&gt;H3</b>	Establishes Metro Mirror relationships between the H2 volumes and the H3 volumes and starts data replication from H2 to H3.
<b>Start H3-&gt;H1</b>	Establishes Metro Mirror relationships between the H3 volumes and the H1 volumes and starts data replication from H3 to H1.
<b>Start H3-&gt;H2</b>	Establishes Metro Mirror relationships between the H3 volumes and the H2 volumes and starts data replication from H3 to H2.
<b>StartGC</b>	<p>Establishes Global Copy relationships between the volumes on the production site and the volumes at the other sites in the session, and begins asynchronous data replication to those sites. For example, if Site 1 is the production site, data replication occurs between the H1 and H2 volumes and the H1 and H3 volumes at the same time.</p> <p>The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.</p>
<b>StartGC H1-&gt;H2</b>	<p>Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2.</p> <p>The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.</p>
<b>StartGC H1-&gt;H3</b>	<p>Establishes Global Copy relationships between the H1 volumes and the H3 volumes, and begins asynchronous data replication from H1 to H3.</p> <p>The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.</p>
<b>StartGC H2-&gt;H1</b>	<p>Establishes Global Copy relationships between the H2 volumes and the H1 volumes, and begins asynchronous data replication from H2 to H1.</p> <p>The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.</p>
<b>StartGC H2-&gt;H3</b>	<p>Establishes Global Copy relationships between the H2 volumes and the H3 volumes, and begins asynchronous data replication from H2 to H3.</p> <p>The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.</p>
<b>StartGC H3-&gt;H1</b>	<p>Establishes Global Copy relationships between the H3 volumes and the H1 volumes, and begins asynchronous data replication from H3 to H1.</p> <p>The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.</p>

Table 104. Metro Mirror multi-target command (continued)

Command	Action
<b>StartGC H3-&gt;H2</b>	Establishes Global Copy relationships between the H3 volumes and the H2 volumes, and begins asynchronous data replication from H3 to H2.  The session remains in the Preparing state until you issue a Start, Suspend, or Terminate command.
<b>Stop</b>	Stops updates to all the target volumes of the pairs in a session.  The following information applies to this and the following <b>Stop</b> commands: <ul style="list-style-type: none"> <li>• Targets are not consistent after the stop.</li> <li>• You can issue the command at any point during an active session.</li> </ul>
<b>StopH1H2</b>	Stops updates from H1 to H2 volumes in a session.
<b>StopH1H3</b>	Stops updates from H1 to H3 volumes in a session.
<b>StopH2H1</b>	Stops updates from H2 to H1 volumes in a session.
<b>StopH2H3</b>	Stops updates from H2 to H3 volumes in a session.
<b>StopH3H1</b>	Stops updates from H3 to H1 volumes in a session.
<b>StopH3H2</b>	Stops updates from H3 to H2 volumes in a session.
<b>Suspend</b>	Suspends updates to all the target volumes of the pairs in a session.  The following information applies to this and the following <b>Suspend</b> commands: <ul style="list-style-type: none"> <li>• You can issue the command at any point during an active session.</li> <li>• Avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.</li> <li>• When a <b>Suspend</b> command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.</li> </ul>
<b>SuspendH1H2</b>	Suspends updates from H1 to H2 volumes in a session.
<b>SuspendH1H3</b>	Suspends updates from H1 to H3 volumes in a session.
<b>SuspendH2H1</b>	Suspends updates from H2 to H1 volumes in a session.
<b>SuspendH2H3</b>	Suspends updates from H2 to H3 volumes in a session.
<b>SuspendH3H1</b>	Suspends updates from H3 to H1 volumes in a session.
<b>SuspendH3H2</b>	Suspends updates from H3 to H2 volumes in a session.

Table 104. Metro Mirror multi-target command (continued)

Command	Action
<b>Terminate</b>	Removes all copy relationships from the hardware during an active session. If you want the targets to be data consistent before you remove their relationship, issue the following commands before you issue the <b>Terminate</b> command: <ul style="list-style-type: none"> <li>• <b>Suspend</b></li> <li>• <b>RecoverHx</b></li> </ul> Where x is the site to which you want to recover.
<b>TerminateH1H2</b>	Removes all H1>H2 or H2>H1 relationships from the hardware during an active session. If you want the target to be data consistent before you remove their relationship, issue the following commands before you issue the <b>Terminate</b> command: <ul style="list-style-type: none"> <li>• <b>SuspendH1H2</b></li> <li>• <b>Recoverx</b></li> </ul> Where x is the site to which you want to recover.
<b>TerminateH1H3</b>	Removes all H1>H3 or H3>H1 relationships from the hardware during an active session. If you want the target to be data consistent before you remove their relationship, issue the following commands before you issue the <b>Terminate</b> command: <ul style="list-style-type: none"> <li>• <b>SuspendH1H3</b></li> <li>• <b>Recoverx</b></li> </ul> Where x is the site to which you want to recover.
<b>TerminateH2H3</b>	Removes all H2>H3 or H3>H2 relationships from the hardware during an active session. <ul style="list-style-type: none"> <li>• <b>SuspendH2H3</b></li> <li>• <b>Recoverx</b></li> </ul> Where x is the site to which you want to recover.

### Example: Setting the production site and starting a Metro Mirror copy to both targets

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Metro Mirror copy to both H1 and H2, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start**

### Example: Setting the production site and starting a Global Copy to both targets

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Global Copy to both H1 and H2, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start GC**



### **Example: Setting the production site and starting a Metro Mirror copy to a single target**

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Metro Mirror copy only to H2, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start H3->H2**

### **Example: Setting the production site and starting Metro Mirror copy to one target and a Global Copy to the second target**

In this example, the production site is site 1 and H1>H2 and H1>H3 are in the Defined state. To change to the production site to site 3 and start a Metro Mirror copy to H2 and a Global Copy to H1, issue the commands in the order that is presented:

1. **Set Production to Site 3**
2. **Start H3->H2**
3. **Start GC H3->H1**

### **Example: Changing the production site after recovering to the site**

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To recover to site 2 and confirm the production site as site 2, issue the

1. **Suspend**
2. **RecoverH2**
3. **Confirm Production at Site 2**

### **Example: Re-enabling the original production site**

In this example, the production site is site 2 and H2>H1 and H2>H3 are active after issuing the commands in the preceding example. To re-enable H1 as the production site, issue the **Re-enable Original Direction** command.

### **Example: Completing a failover to a single target for practice**

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To failover to site 3, issue the following commands in the order that is presented:

1. **SuspendH1H3**
2. **FailoverH3**

If the production site is site 2 and H2>H1 and H2>H3 are active, issue the following commands to failover to site 3:

1. **SuspendH2H3**
2. **FailoverH3**

### **Example: Suspending the copy of data to both target sites**

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To suspend the copy of data to both target sites, issue the **Suspend** command.

### **Example: Suspending the copy of data to a single target site**

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To suspend the copy of data only from H1 to H3, issue the **SuspendH1H3** command.

### **Example: Stopping the copy of data to both target sites**

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To stop the copy of data to both target sites, issue the **Stop** command.

### **Example: Stopping the copy of data to a single target site**

In this example, the production site is site 1 and H1>H2 and H1>H3 are active. To stop the copy of data only from H1 to H3, issue the **StopH1H3** command.

## **Site awareness**

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

**Note:** To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A\_Location and a session with the location of A\_Location for its H1 role, changing the location of the storage system to a different location, such as B\_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A\_Location, the session's role location is not changed.

**Important:** Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

## Preserve Mirror option

This topic presents recommendations for using the Preserve Mirror option in FlashCopy and Metro Mirror sessions.

When the source of the FlashCopy relationship is a source of a Metro Mirror relationship, and the target of the FlashCopy relationship is the source of a Metro Mirror relationship, the Preserve Mirror option attempts to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship, preventing a full copy from being performed over the Metro Mirror link. Instead, parallel flashes will be performed (if possible) on both sites. If the consistency cannot be preserved, the Flash for the FlashCopy relationships will fail, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship will not be changed.

**Note:** This option is available only on DS8000 storage devices with the required code levels installed.

However, in some instances, the Preserve Mirror option can cause a Metro Mirror session to go into a Preparing state, or even a Suspended state. This topic describes the recommended usage of the Preserve Mirror feature. Using this feature in other ways might lead to a Metro Mirror session going into a Preparing or Suspended state.

### FlashCopy session

You can use the Preserve Mirror option in FlashCopy sessions in two different methods:

#### Perform an incremental resynchronization

To perform an incremental resynchronization, select the **Incremental** and **Persistent** options in the FlashCopy session: do *not* select the **No Copy** option.

#### Perform a single full copy

To perform a single full copy, ensure that the **Incremental**, **Persistent** and **No Copy** options are *not* selected before you issue a **Flash** command. If you use the **No Copy** option, issue either an **Initiate Background Copy** command or **Terminate** command before you issue the Flash command.

Refer to your DS8000 documentation for more information about the Preserve Mirror function.

### Metro Mirror session

You can set up your Metro Mirror pairs in two different ways, depending on the level of consistency you need, and your preferences.

**Note:** For the examples in this section, the source pair is H1a->H2a and the target pair is H1b->H2b. The source pair will contain volumes that will be the source of the FlashCopy relationship and the target pair will always contain volumes that will be the target of the FlashCopy relationship.

#### Create one Metro Mirror session, and add the Metro Mirror pairs as copy sets to that session

The benefit to this approach is that you do not need to worry about whether the host considers the H1a->H2a and H1b->H2b volumes to be

consistent with one another. IBM Tivoli Storage Productivity Center for Replication will ensure that all of the volumes remain consistent.

A drawback to this approach is that when using the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required), there is a chance that the target pair (H1b->H2b) might suspend unexpectedly: this causes all other pairs in the Metro Mirror session to suspend (including H1a->H2a). This can occur when a FlashCopy establish or withdraw fails unexpectedly on the remote (H1b->H2b) site. If the host requires the H1a->H2a and H1b->H2b volumes to be consistent, then you should suspend all other volumes.

#### **Create one Metro Mirror session for the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes**

Use this option when the hosts and applications do not require the H1a->H2a and H1b->H2b volumes to be consistent with one another. In this case, you should create one Metro Mirror session for all of the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes. The H1a->H2a pair is added to the first session, while the H1b->H2b pair is added to the second Metro Mirror session. As long as the host does not require consistency between the H1a and H1b volumes, this option benefits you when you use the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required). The benefit is that if one pair is suspended (such as H1a->H2a), the pairs in the other session will not be affected, since it is in a different Metro Mirror session. Using this method, you can avoid the situation in which a critical application is writing to the source pair (H1a->H2a), while a batch job is writing to the target pair (H1b->H2b), and both pairs are in the same IBM Tivoli Storage Productivity Center for Replication session. These factors cause both applications to receive extended long busy signals, instead of just the batch job.

## **Creating sessions and adding copy sets**

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

### **Creating a FlashCopy session and adding copy sets**

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy session was established.

When you create a FlashCopy session for Global Mirror or Metro Global Mirror space-efficient target volumes, you must select **No Copy** for the FlashCopy session. With space-efficient volumes, you can use your FlashCopy repository more efficiently. Instead of requiring an equal amount of space to write data to, you can set aside a smaller amount of space in which to write data, where only the tracks that are changed are recorded. When your pool of storage is full, you can no longer perform a FlashCopy operation, and your session goes into a Severe state.

To add a FlashCopy session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a FlashCopy session:
  - a. In the menu bar, click **Sessions**.
  - b. On the Sessions page, click **Create Session**.
  - c. In the Create Session window, select the following options:

**Hardware type**

Select the type of storage system for the session:

- DS8000, DS6000, ESS 800
- SAN Volume Controller
- Storwize Family (for storage systems in the IBM Storwize family)
- Storwize V7000 Unified

**Session type**

Select **FlashCopy**.

**Session name**

Enter a name to identify the session. The session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

**Site 1 Location**

For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

**Session image**

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session. The lightning bolt in the image represents FlashCopy replication.

If you selected a location for the site that is in the session, the site is labeled with the site location. If you did not select site locations, the sites are labeled with **Site x**, where *x* is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

- d. Click **OK**.
  - e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.
2. Follow these steps to add copy sets to the session:
  - a. On the Select Host1 page of the Add Copy Sets wizard, complete the following information. The field names that are displayed depend on the storage system type. When you complete the information, click **Next**.

**Storage system**

Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the

location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

**Logical storage system or I/O Group**

Select a logical subsystem (LSS) or I/O group.

**Volume**

Select one volume or all volumes. The volumes are limited to the volumes within the LSS or I/O group that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

**Session image**

Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

**Volume Details**

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

**Use a CSV file to import copy sets**

Select this option to import copy sets from a comma-separated value (CSV) file. Click **Browse** to select the CSV file.

- b. On the Choose Target1 page, select the target storage system, LSS or I/O group, and volume. Click **Next**.
- c. On the Select Copy Sets page, select from the following options and click **Next**:

**Select All**

Click this button to select all of the copy sets in the table.

**Deselect All**

Click this button to clear all of the copy sets in the table.

**Add More**

Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. On this page, the **Storage system** and **Logical storage system or I/O Group** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the Choose Host1 and Choose Target1 pages and add additional copy sets.

**Selection check boxes**

Select one or more copy sets that you want to create.

**Host 1** Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

**Copy Set**

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
  - e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.
3. Follow these steps to add properties for the session:
- a. On the Sessions page, select the session.
  - b. From the **Session Actions** list, select **View/Modify > Properties**.
  - c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see “Viewing session properties” on page 447.
  - d. Click **OK**.

**Related concepts:**

“Sessions” on page 340

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

“Session types” on page 57

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

## Creating a Snapshot session and adding copy sets

A Snapshot session is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. Snapshot sessions are available only for the XIV system.

To add a Snapshot session, you create the session and then add copy sets to the session.

1. Follow these steps to create a Snapshot session:
  - a. In the menu bar, click **Sessions**.
  - b. On the Sessions page, click **Create Session**.
  - c. In the Create Session window, select the following options:

**Hardware type**

Select **XIV**.

**Session type**

Select **Snapshot**.

**Session name**

Enter a name to identify the session. The session name can be up to 58 alphanumeric characters. The default name for a new session is **newsession**.

**Site 1 Location**

For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that only storage systems that have either a matching site location for each role or have no defined location are displayed in the copy set wizard.

### Session image

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

If you selected a location for the site that is in the session, the site is labeled with the site location. If you did not select site locations, the sites are labeled with **Site x**, where *x* is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

- d. Click **OK**.
  - e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.
2. Follow these steps to add copy sets to the session:
- a. From the **Host1 storage system** list, select the storage system that contains the volumes that you want to add.  
  
If the H1 role has an assigned location, only those storage systems that have the same location as the H1 role or storage systems that do not have a set location are displayed for selection.  
  
If the H1 role does not have an assigned location, all storage systems are displayed for selection.  
  
Storage systems that are assigned to a location are listed under the location name. Storage systems that are not assigned to a location are listed under the **None**. column.
  - b. From the **Host1 pool** list, select the pool that contains the volumes.
  - c. From the **Host1 volume** list, select the volumes. To select multiple volumes, press Ctrl or Shift and click the volumes in the list.
  - d. If you want to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Click **Browse** to select the CSV file and click **Next**.
  - e. On the Matching Results page, click **Next** if the match was successful.
  - f. On the Select Copy Sets page, select from the following options and click **Next**.

### Select All

Click this button to select all of the copy sets in the table.

### Deselect All

Click this button to clear all of the copy sets in the table.

### Add More

Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. The lists on this page are populated with the values from the previously selected copy set. Repeat the previous steps to complete the Choose Host1 page and add additional copy sets.



### Selection check boxes

Select one or more copy sets that you want to create.

**Host 1** Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

### Copy Set

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- g. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- h. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

If you want to add a description for the session, select the session on the Sessions page. From the **Session Actions** list, select **View/Modify > Properties**.

### Related concepts:

“Sessions” on page 340

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

“Session types” on page 57

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

## Creating a Metro Mirror session and adding copy sets

A Metro Mirror session is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart.

To add a Metro Mirror session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a Metro Mirror session:
  - a. In the menu bar, click **Sessions**.
  - b. On the Sessions page, click **Create Session**.
  - c. In the Create Session window, select the following options:

### Hardware type

Select the type of storage system for the session:

- DS8000, DS6000, ESS 800
- SAN Volume Controller
- Storwize Family (for storage systems in the IBM Storwize family)
- Storwize V7000 Unified
- XIV

### Session type

Select a Metro Mirror session type.

### Session name

Enter a name to identify the session. For sessions that contain an IBM XIV Storage System, the session name can be up to 58 alphanumeric characters. For sessions that contain other storage system types, the session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

### Site *x* Location

For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

### Session image

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

This image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:

- Straight arrows = synchronous replication
- Straight arrows with an asynchronous replication symbol = asynchronous replication
- Lightning bolts = FlashCopy replication

If you selected a location for the sites that are in the session, the sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site *x***, where *x* is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

- d. Click **OK**.
  - e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.
2. Follow these steps to add copy sets to the session:
    - a. In the Add Copy Sets wizard, enter the following information for the choose volume pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

### Storage system

Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

### Logical storage system, I/O Group, or Pool

Select a logical subsystem (LSS), I/O group, or pool.

### Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

### Session image

Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

### Volume Details

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

### Use a CSV file to import copy sets

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- b. On the Select Copy Sets page, select from the following options and click **Next**:

#### Select All

Click this button to select all of the copy sets in the table.

#### Deselect All

Click this button to clear all of the copy sets in the table.

#### Add More

Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. On this page, the **Storage system** and **Logical storage system, I/O Group, or Pool** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the choose volume pages and add additional copy sets.

### Selection check boxes

Select one or more copy sets that you want to create.

- Host 1** Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

### Copy Set

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- c. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- d. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

3. Follow these steps to add properties for the session:
  - a. On the Sessions page, select the session.
  - b. From the **Session Actions** list, select **View/Modify > Properties**.
  - c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see “Viewing session properties” on page 447.
  - d. Click **OK**.

**Related concepts:**

“Sessions” on page 340

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

“Session types” on page 57

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

### **Creating a Global Mirror session and adding copy sets**

A Global Mirror session is a method of asynchronous, remote data replication between two sites that are over 300 KM apart.

To add a Global Mirror session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a Global Mirror session:
  - a. In the menu bar, click **Sessions**.
  - b. On the Sessions page, click **Create Session**.
  - c. In the Create Session window, select the following options:

**Hardware type**

Select the type of storage system for the session:

- DS8000, DS6000, ESS 800
- SAN Volume Controller
- Storwize Family (for storage systems in the IBM Storwize family)
- Storwize V7000 Unified
- XIV

**Session type**

Select a Global Mirror session type.

**Session name**

Enter a name to identify the session. For sessions that contain an IBM XIV Storage System, the session name can be up to 58 alphanumeric characters. For sessions that contain other storage system types, the session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

**Site *x* Location**

For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only

storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

### Session image

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

This image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:

- Straight arrows = synchronous replication
- Straight arrows with an asynchronous replication symbol = asynchronous replication
- Lightning bolts = FlashCopy replication

If you selected a location for the sites that are in the session, the sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site x**, where *x* is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

- d. Click **OK**.
  - e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.
2. Follow these steps to add copy sets to the session:
    - a. In the Add Copy Sets wizard, enter the following information for the choose volume pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

### Storage system

Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

### Logical storage system, I/O Group, or Pool

Select a logical subsystem (LSS), I/O group, or pool.

### Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

### Session image

Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

### Volume Details

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

### Use a CSV file to import copy sets

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- b. On the Select Copy Sets page, select from the following options and click **Next**:

#### Select All

Click this button to select all of the copy sets in the table.

#### Deselect All

Click this button to clear all of the copy sets in the table.

#### Add More

Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. On this page, the **Storage system and Logical storage system, I/O Group, or Pool** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the choose volume pages and add additional copy sets.

#### Selection check boxes

Select one or more copy sets that you want to create.

**Host 1** Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

#### Copy Set

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- c. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
  - d. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.
3. Follow these steps to add properties for the session:
    - a. On the Sessions page, select the session.
    - b. From the **Session Actions** list, select **View/Modify > Properties**.
    - c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see “Viewing session properties” on page 447.
    - d. Click **OK**.

#### Related concepts:

“Sessions” on page 340

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

“Session types” on page 57

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

## Creating a Metro Global Mirror session and adding copy sets

A Metro Global Mirror session is a method of continuous, remote data replication that operates between three sites of varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

To add a Metro Global Mirror session, you create the session, add copy sets to the session, and then select the properties for the session.

1. Follow these steps to create a Metro Global Mirror session:
  - a. In the menu bar, click **Sessions**.
  - b. On the Sessions page, click **Create Session**.
  - c. In the Create Session window, select the following options:

### Hardware type

Select the hardware type **DS8000, DS6000, ESS 800**.

### Session type

Select a Metro Global Mirror session type.

### Session name

Enter a name to identify the session. The session name can be up to 250 alphanumeric characters. The default name for a new session is `newsession`.

### Site *x* Location

For each site that is in the session, select the location for the site or **None**. The site location, if available, is set by storage system on the Storage Systems page.

You can associate a location with each storage system and each site in a session. This site awareness ensures that you can select only storage systems that have either a matching site location for each role or have no defined location when you add copy sets to a session.

### Session image

Shows an image that represents the session. The session image is displayed when you select a session type. The image is a visual aid to help you create your session.

This image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:

- Straight arrows = synchronous replication
- Straight arrows with an asynchronous replication symbol = asynchronous replication
- Lightning bolts = FlashCopy replication

If you selected a location for the sites that are in the session, the sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site x**, where *x* is the site number. You can hover the mouse pointer over the site location to view the text in a larger format.

- d. Click **OK**.
  - e. If the session was successfully created, click **Launch Add Copy Sets Wizard** and continue to the next step. If the session was not created, refer to the displayed message text to fix the problem.
2. Follow these steps to add copy sets to the session:
- a. In the Add Copy Sets wizard, enter the following information for the choose volume pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

#### **Storage system**

Select a storage system. If the volume role has a site location that is assigned to it, you can select a storage system that is assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the location. Storage systems that do not have a location are listed under **None**. If the role does not have a site location, you can select any storage system.

#### **Logical storage system**

Select a logical subsystem (LSS).

#### **Volume**

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system.

#### **Session image**

Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

#### **Volume Details**

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

#### **Use a CSV file to import copy sets**

Select this option to import copy sets from a comma-separated value (CSV) file. Click **Browse** to select the CSV file.

- b. On the Select Copy Sets page, select from the following options and click **Next**:

#### **Select All**

Click this button to select all of the copy sets in the table.

#### **Deselect All**

Click this button to clear all of the copy sets in the table.

#### **Add More**

Click this button to add another copy set to the list of copy sets to be created.



When you click **Add More**, you are returned to the Choose Host1 page of the wizard. On this page, the **Storage system** and **Logical storage system, I/O Group, or Pool** lists are populated with the values from the previously selected copy set. Repeat the previous steps to complete the choose volume pages and add additional copy sets.

### Selection check boxes

Select one or more copy sets that you want to create.

**Host 1** Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

### Copy Set

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- c. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
  - d. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.
3. Follow these steps to add properties for the session:
- a. On the Sessions page, select the session.
  - b. From the **Session Actions** list, select **View/Modify > Properties**.
  - c. In the View/Modify Properties notebook, select the options that you want. For a description of the options, see “Viewing session properties” on page 447.
  - d. Click **OK**.

### Related concepts:

“Sessions” on page 340

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

“Session types” on page 57

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

## Completing session administration tasks

Use session commands to complete session tasks such as starting, stopping, suspending, or terminating sessions.

You can issue session commands by using the GUI, as described in this topic, or you can use the CLI to issue commands. For a description of the CLI session commands, see the IBM Tivoli Storage Productivity Center Command-line Interface Reference.

To issue a command for a session, complete the following steps:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to issue a command.
3. Click **Session Actions > Commands**.

If the session is a single-target session type, all available commands are displayed under **Commands**. Click the command that you want to issue.

If the session is a multi-target session type, some of the commands are grouped by overall function. For example, **Start**, **Stop**, or **Suspend**. When you click a command, a window is displayed that provides specific commands for that function. For example, for Metro Mirror - Metro Mirror sessions, if you select **Start**, a window is displayed that contains the specific start commands that you can select. Select the command that you want, review the message text for that command, and then click **Yes** if you want the command to run.

**Related reference:**

“Session commands” on page 383

The commands that are available for a session depend on the session type, and for some commands, the storage system type.

## Using the Metro Mirror heartbeat

This topic provides information about Metro Mirror heartbeat, including how to enable and disable the heartbeat.

### Metro Mirror heartbeat

The *heartbeat* is a Metro Mirror function. When the Metro Mirror heartbeat is disabled, data consistency across multiple storage systems is not guaranteed if the Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The problem occurs as a result of the Hardware Freeze Timeout Timer function within the storage system. If the controlling software loses connection to a storage system, the Metro Mirror relationships that it is controlling remains established and there is no way to *freeze* those pairs to create consistency across the multiple storage systems. When the freeze times out, dependent I/O is written to the target storage systems, which might corrupt data consistency. Freeze refers to a Metro Mirror (peer-to-peer remote copy [PPRC]) freeze function.

When determining whether to use the Metro Mirror heartbeat, analyze your business needs. Disabling the Metro Mirror heartbeat might result in data inconsistency. If you enable the Metro Mirror heartbeat and a freeze occurs, your applications will be unable to write during the freeze.

Metro Mirror heartbeat is disabled by default.

Metro Mirror heartbeat is not available for Metro Mirror with HyperSwap or Metro Global Mirror with HyperSwap.

There are two cases where lost communication between the coordination software (controller) and one or more storage systems can result in data consistency loss:

#### Freeze event not detected by a disconnected storage system

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to prevent any writes from occurring, for a period determined by the Freeze timeout timer. At the same time, the affected

storage controller loses communication with the controlling software and cannot communicate the Freeze event to the software.

Unaware of the problem, the controlling software does not issue the Freeze command to the remaining source storage systems. The freeze will stop dependent writes from being written to connected storage systems. However, once the Freeze times out and the long-busy is terminated, dependent write I/Os continue to be copied from the storage systems that did not receive the Freeze command. The Metro Mirror session remains in a state where one storage system has suspended copying while the other three storage systems are still copying data. This state causes inconsistent data on the target storage systems.

#### **Freeze event detected, but unable to propagate the Freeze command to all storage systems**

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to issue long-busy to the applications for a period determined by the Freeze timeout timer. At the same time, one of the remaining three source systems loses communications with the controlling software.

The storage system that had an error writing to its target cannot communicate the Freeze event to the controlling software. The controlling software issues the Freeze command to all but the disconnected storage system (the one that lost communication with the software). The long-busy stops dependent writes from being written to the connected storage systems.

However, once the Freeze times out on the frozen storage system and the long-busy is terminated, dependent write I/Os continue to the target storage system from the source storage system that lost communication and did not receive the Freeze command. The Metro Mirror session remains in a state where three storage systems have suspended copying and one storage system is still copying data. This state causes inconsistent data on the target storage systems.

Before Tivoli Storage Productivity Center for Replication V3.1, if the controlling software within a Metro Mirror environment detected that a managed storage system lost its connection to its target, the controlling software stopped all the other source systems to ensure consistency across all the targets. However, if the controlling software lost communication with any of the source subsystems during the failure, it could not notify those storage systems of the freeze event or ensure data consistency. The Metro Mirror heartbeat helps to overcome this problem. In a high-availability configuration, the Metro Mirror heartbeat is continued by the standby server after the Takeover command is issued on the standby, enabling you to perform actions on the standby server without causing a freeze.

Tivoli Storage Productivity Center for Replication registers with the managed ESS 800, DS6000 or DS8000 storage systems within a Metro Mirror session when the start command is issued to the session. After this registration occurs, a constant heartbeat is sent to the storage system. If the storage system does not receive a heartbeat from the Tivoli Storage Productivity Center for Replication management server within the allotted time (a subset of lowest LSS timeout value across all the source LSSs), the storage system initiates a freeze. If Tivoli Storage Productivity Center for Replication did not successfully communicate with the storage system, it initiates a freeze on the remaining storage system after the allotted time is expired.

**Note:** It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.

When you are using the Metro Mirror heartbeat, be aware that:

- The Metro Mirror heartbeat can cause a single point of failure: if an error occurs on just the management server and not the storage system, a freeze might occur.
- When the Metro Mirror heartbeat timeout occurs, the storage system remains in a long busy state for the duration of the LSS freeze timeout.

**Note:** If Metro Mirror heartbeat is enabled for storage systems that are connected through a HMC connection, a connection loss might cause lost heartbeats, resulting in Freeze actions with application I/O impact for configured Extended Long Busy timeout.

The Metro Mirror heartbeat is supported on storage systems connected through a TCP/IP (direct connect or HMC) connection. It is not supported on storage systems connected through a z/OS connection. Enabling the Metro Mirror heartbeat with a z/OS connection does not fail; however, a warning message is displayed specifying that the Metro Mirror heartbeat function does not work unless you have an IP connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP (either direct connect or HMC) connection and z/OS connection, and the TCP/IP connection fails, Tivoli Storage Productivity Center for Replication suspends the Metro Mirror session because there is no heartbeat through the z/OS connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP connection and z/OS connection and you remove all TCP/IP connections, Tivoli Storage Productivity Center for Replication suspends the Metro Mirror sessions and the applications using those volume will be in Extended Long Busy timeout until the storage system's internal timeout timer expires. Ensure that you disable the Metro Mirror heartbeat for all Metro Mirror sessions before removing the last TCP/IP connection to avoid the Extended Long Busy timeout.

### **Enabling and disabling the Metro Mirror heartbeat**

The Metro Mirror heartbeat guarantees data consistency across multiple storage systems when the Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The Metro Mirror heartbeat is disabled by default.

To enable the Metro Mirror heartbeat, perform the following steps:

1. In the menu bar, click **Settings > Advanced Tools**.
2. To enable the Metro Mirror heartbeat, click **Enable Heartbeat**.
3. To disable the Metro Mirror heartbeat, click **Disable Heartbeat**.

## **Exporting copy set data**

You can export data about all copy sets in a specific session to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Complete these steps to export the copy sets in a specific session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to export copy sets.
3. From the **Session Actions** list, select **Export > Export Copy Sets**. The Export Copy Sets window displays the status of the export and a link to the exported file if the export completed.
4. Right-click the link and save the file to a local system.

## Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session for which you want to import copy sets.
3. From the **Session Actions** list, select **View/Modify > Add Copy Sets**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Click **Browse**, select the CSV file, and then click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets that you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and then click **Next**.
9. View the information on the Results page of the wizard, and then click **Finish**.

## Exporting historical data for Global Mirror role pairs

You can export data for a Global Mirror role pair that is in a TotalStorage Enterprise Storage Server, System Storage DS6000, or System Storage DS8000 session to a comma-separated value (CSV) file. You can then use the data in the CSV file to analyze trends in your storage environment that affect your recovery point objective (RPO).

Complete these steps to export data to a CSV file:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, click the Global Mirror session for which you want to export data.
3. From the **Session Actions** list, select **Export > Export Global Mirror Data**.
4. In the Export Historical Data for Global Mirror window, complete the following information, and then click **OK**:

### Select the role pair

Select the role pair for the data that you want to show in the CSV file.

### Select the type of data

Select the type of data that you want to export to a CSV file.

Depending on the data type that you select, the CSV file contains data about the RPO or data about logical subsystem (LSS) out-of-sync tracks.

To better analyze trends, you can create a CSV file with one data type and then create a file with the other data type. For example, the file that contains data for the RPO might show that the RPO threshold is

often exceeded on a particular day and time. You can then view the file that contains data for LSS out-of-sync tracks to see whether a particular LSS or set of LSSs have high out-of-sync track values for that day and time.

**RPO** Select this option to export data for the RPO. The data that is shown in the export file includes the average RPO for the dates that you select and information related to the formation of consistency groups.

**LSS Out-of-Sync Tracks**

Select this option to export data for the out-of-sync tracks that are in the LSSs.

**Start date**

Select the start date for the data that you want to show in the export file.

By default, the date range maximum for RPO files is 31 days of data and the maximum for LSS out-of-sync track files is 7 days of data.

**End date**

Select the end date for the data that you want to show in the export file.

After you click **OK** to export the file, a results window displays the status of the export and a link to the exported file if the export completed. Download and save the file to a local system.

## Modifying the location of session sites

You can change the location that is associated with each site in session.

**Prerequisites:** You must have Administrator privileges or Operator privileges for the session to modify the location of a site in a session.

Changing the location of a site in a session does not affect the location of the storage systems that are associated with that site.

Complete the following steps to modify the location of a site:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session whose site locations you want to change.
3. From the **Session Actions** list, select **View/Modify > Site Location(s)**.
4. In the Modify Site Locations window, select a site location or **None** for each site that you want to change. An image that represents the sites in the session is displayed. If you selected locations for the sites that are in the session, the sites are labeled with the site location. If you did not select site locations, the sites are labeled with **Site *n***, where *n* is the site number. You can hover the mouse pointer over the site name to view the name in a larger format.
5. Click **OK**.

**Related concepts:**

“Site awareness” on page 323

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

## Removing sessions

This topic describes how to remove sessions.

**Important:** You can remove only sessions that are in the Defined state.

Complete these steps to remove a session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session that you want to remove.
3. From the **Session Actions** list, select **Remove Session**.
4. Click **Yes** to confirm that you want to remove the session.

## Removing copy sets

This topic describes how to remove copy sets.

Complete these steps to remove a copy set:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session that you want to remove copy sets from.
3. From the **Session Actions** list, select **View/Modify > Remove Copy Sets**. This starts the Remove Copy Sets wizard.
4. From the drop-down menus in the Remove Copy Sets wizard, select the Host 1 storage system, logical storage subsystem, and volume or select the all option. If you select all for a filter, the lower-level filter or filters are disabled. Click **Next**.
5. Select the copy sets that you want to remove and click **Next**.
6. The number of copy sets to be removed is displayed. Select the following options for removing the copy sets and click **Next**:
  - **Do you want to keep the base relationships on the hardware, but remove the copy sets from the session?**
    - **Yes.** This option specifies that the base relationships remain on the hardware, but the copy sets are removed from the Tivoli Storage Productivity Center for Replication session. This option supports scenarios in which it might be best to leave the relationship on the hardware to avoid performing a full copy. For example, when you are migrating from one session type to another.  
Only the base relationships (Metro Mirror, Global Copy, Snapshot, and FlashCopy) remain on the hardware. The relationships are removed from any consistency groups that are defined on the storage system.
    - **No.** This option specifies that all relationships for the copy sets are removed from the hardware as well as the Tivoli Storage Productivity Center for Replication session. This option is the default.
  - **If there are errors removing relationships on the hardware, do you want to force the copy sets to be removed from the session?**
    - **Yes.** This option forces the removal of copy sets despite any errors that occur when removing the relationships from the storage system. Once a forced removal is complete, any relationships that remain on the storage system for that copy set must be removed manually using the storage system interface.
    - **No.** This option does not force the removal of copy sets. This option enables you to correct the errors and try to remove the copy sets again. This option is the default.

7. After the copy sets are removed, click **Finish**.

**Related concepts:**


“Copy sets” on page 341

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

## Migrating existing hardware relationships to Tivoli Storage Productivity Center for Replication

You can convert existing hardware relationships for copy services to a Tivoli Storage Productivity Center for Replication session.

For session types other than multi-target sessions, you can either complete this action manually as described in this topic or use the data migration utility for Tivoli Storage Productivity Center for Replication. The data migration utility is not available for multi-target sessions.

To download the data migration utility, go to  <http://www.ibm.com/support/docview.wss?uid=swg24000625>. The data migration utility produces the Tivoli Storage Productivity Center for Replication CLI command script files and DSCLI script files that necessary to migrate the configuration to Tivoli Storage Productivity Center for Replication.

To manually migrate role pairs in a hardware relationship, complete the following steps:

1. Identify the hardware relationships that you want to migrate.
2. Create a Tivoli Storage Productivity Center for Replication session and add copy sets to the session.
3. If you are migrating from one of the following storage systems, complete the applicable step:
  - a. For DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems with Global Mirror relationships, terminate the Global Mirror master.
  - b. For SAN Volume Controller, Storwize, or the XIV system storage systems, remove all volumes from the consistency group.
4. Issue a start command to start the Tivoli Storage Productivity Center for Replication session.

**Related reference:**

“Session commands” on page 383

The commands that are available for a session depend on the session type, and for some commands, the storage system type.

---

## Practicing disaster recovery

You can use practice volumes to test your disaster recovery actions while maintaining disaster recovery capability. Practice volumes are available in Metro Mirror Failover and Failback sessions, Global Mirror Failover and Failback, Global Mirror Either Direction sessions, and Metro Global Mirror with Practice sessions.

**Important:** You can test your disaster recovery actions without using practice volumes. However, without practice volumes, you cannot maintain disaster recovery capability while continuing to copy the data.



## Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

**Note:** You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

## Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session

A Metro Mirror Failover and Failback session with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This function is available on the following storage systems:

- System Storage DS8000
- System Storage DS6000
- TotalStorage Enterprise Storage Server Model 800
- SAN Volume Controller,
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

Perform these steps to practice disaster recover actions for a Metro Mirror Failover/Failback with Practice session:

1. Start a Metro Mirror with Practice session.
2. When the Metro Mirror session reaches the Prepared state, issue a Flash command to make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

**Note:** For ESS, DS6000, and DS8000 storage systems, the **Flash** command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

## Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session

A Global Mirror (either direction) with two-site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

**Note:** This function is available only on ESS, DS6000, and DS8000 storage systems.

Perform these steps to practice disaster recover actions for a Global Mirror either Direction with two-site Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

**Note:** FlashCopy must be always a full copy due to limitations of the hardware.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

**Note:** With two directions, you can reverse the direction of your data flow.

## Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session

A Global Mirror Failover and Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred.

You can do this practice without losing your disaster recovery capability. The number of volumes used for the device varies, but the steps to conduct a Global Mirror Failover and Failback with Practice are the same for both devices.

This function is available on the following storage systems:

- System Storage DS8000
- System Storage DS6000
- TotalStorage Enterprise Storage Server Model 800
- SAN Volume Controller,
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

Perform these steps to practice disaster recover actions for a Global Mirror Failover/Failback with Practice session:

1. Start a Global Mirror with Practice session.

2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

**Note:** FlashCopy must be always a full copy due to limitations of the hardware.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

## Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session

A Metro Global Mirror Failover/Failback with Practice session combines Metro Mirror, Global Mirror and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this to practice what you might do if a disaster occurred without losing your disaster recovery capability.

**Note:** This function is available on ESS, DS6000, and DS8000 storage systems.

The intermediate volume is on the third site (I3). This maintains disaster recovery capability while a copy is kept on the H3 volume for practice purposes.

Perform these steps to practice disaster recover actions for a Metro Global Mirror Failover/Failback with Practice session:

1. Start a Metro Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to take a point-in-time copy of the data that is on I3, on H3. This creates a consistent point-in-time copy of your data on the H3 volume, and then automatically restarts the session so that copying from H1 to H2 to I3 continues. The Flash command temporarily stops copying the data from site 2 to site 3, in order to create a consistent point-in-time copy on I3, while maintaining disaster recovery capabilities on site 2 using the Metro Mirror portion of the session. Then, data replication from H2 to I3 is restarted.

**Note:** FlashCopy must be always a full copy due to hardware limitations.

3. Practice the same actions you would take in an actual disaster, using H3 as your practice host volume, while you run the real application on H1. This enables you to use the same scripts and commands to run on H3 that you would use in an actual disaster.

---

## Monitoring health and status

### Viewing the health summary

Use the Overview page to view overall health and status of sessions, storage systems, host systems, and management servers.

The Overview page is the first page that you see after you log on. You can display this page by selecting **Overview** in the Tivoli Storage Productivity Center for Replication menu bar. This panel provides the following information:

**Overall session status**

Indicates session status, which can be normal, warning, or severe. The status can also be inactive, if all sessions are defined or if no sessions exist.

**Overall storage system status**

Indicates the connection status of storage systems.

**Overall host system status**

Indicates the connection status of host systems.

**Management server status**

Indicates the status of the standby server if you are logged on to the local server. If you are logged on to the standby server, this status indicates the status of the local server.

## Viewing SNMP alerts

You can view Tivoli Storage Productivity Center for Replication SNMP trap descriptions from the Alert Log node in the Tivoli Storage Productivity Center stand-alone GUI.

In the stand-alone GUI navigation tree, expand **Alerting > Alert Log** and click **Replication**.

## Viewing sessions

The Tivoli Storage Productivity Center for Replication GUI includes icons, images, messages, and other visual and informational aids to help you to determine the status and state of your sessions.

**Customizing the session list**

The Session page provides a list of created sessions. This list is organized in columns and rows. There is a row for each session and columns for information that is related to that session, such as the session name, status, state, and type. You can customize the session list by filtering information, resizing columns, and sorting rows.

**Filtering the session list**

Use filtering when you want to limit the rows that are shown in the session list based on the session name or type. To apply a filter, enter the filter text in the **Filter** field on the Session page. You can enter complete text or partial text. Only those rows that contain matching text in the **Name** or **Type** column are displayed in the list.

To remove the filtering for the list, click the **X** icon in the **Filter** field.

**Resizing the columns in the session list**

You can resize columns to more easily view the information in the list. To resize a column, drag the column heading border until the column is at the width that you want.

**Sorting the session list**

By default, the sessions in the list are sorted by session name in ascending order. You can change the value that is used for sorting by clicking the appropriate column heading. For example, if you want to sort by the session type, click the **Type** heading. The list sorts by the value in that column. You can choose an





ascending or descending sort order by clicking the column heading.

## Session status icons

The IBM Tivoli Storage Productivity Center for Replication GUI uses icons to represent the status of each session.

The following table describes each session status icon.

Table 105. Session status icons

Icon	Meaning	Description
	Inactive	The session is in a defined state, with no activity on the hardware.
	Normal	A consistent copy of the data either exists or is being maintained.
	Warning	For Metro Mirror, Global Mirror, and Metro Global Mirror, the session might have volumes that are being synchronized or are about to be synchronized, with no suspended volumes. For FlashCopy, the warning status is valid only after the start command is issued and before the flash. This warning status means that the session is either preparing or is ready for a flash command but targets do not yet have a consistent copy.  If a HyperSwap session is degraded, which means it is enabled on one or more sysplex members and disabled on at least one sysplex member, then the session is in a warning state.
	Severe	One or more errors must be dealt with immediately. Possible causes include the following: <ul style="list-style-type: none"> <li>• One or more volumes are suspended</li> <li>• A session is suspended</li> <li>• A volume is not copying correctly</li> </ul>

## Session images

The Tivoli Storage Productivity Center for Replication GUI provides a visual aid to help you create and manage your sessions. The visual aid shows the number of volume roles in the session and how the roles are distributed between the sites. It also shows the copy method and direction.

## Volume role symbols

The volume role symbols represent the replication status on the volumes.

Table 106. Volume role symbols







Symbol	Description	Meaning
	Active host volumes	This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write input/output (I/O).

Table 106. Volume role symbols (continued)

Symbol	Description	Meaning
	Active host volumes with change volumes	This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write I/O and change volumes.
	Recoverable volumes	This symbol represents volumes that contain a consistent copy of the data.
	Recoverable volumes with change volumes	This symbol represents volumes and change volumes that contain a consistent copy of the data.
	Inconsistent volumes	This symbol represents the volumes that do not contain a consistent copy of the data.
	Inconsistent volumes with change volumes	This symbol represents the volumes and change volumes that do not contain a consistent copy of the data.

### Data copying symbols

The data copying symbols indicate the type of copy that occurs between the volume roles. The direction that the symbol is displayed in the Tivoli Storage Productivity Center for Replication GUI depends on the direction of the copy.

Table 107. Data copying symbols
















Symbol	Type of Copy	Meaning
	FlashCopy	This symbol represents a FlashCopy operation.

Table 107. Data copying symbols (continued)

Symbol	Type of Copy	Meaning
	FlashCopy with errors	This symbol represents a FlashCopy operation with errors on one or more pair.
	FlashCopy inactive	This symbol represents an inactive FlashCopy operation.
	FlashCopy inactive with errors	This symbol represents an inactive FlashCopy operation with errors on one or more pair.
	Synchronous	This symbol represents a synchronous copy.
	Synchronous with errors	This symbol represents a synchronous copy with errors on one or more pairs.
	Synchronous inactive	This symbol represents an inactive synchronous copy.
	Synchronous inactive with errors	This symbol represents an inactive synchronous copy with errors on one or more pair.
	Asynchronous	This symbol represents an asynchronous copy.
	Asynchronous with errors	This symbol represents an asynchronous copy with errors on one or more pair.
	Asynchronous inactive	This symbol represents an inactive asynchronous copy.
	Asynchronous inactive with errors	This symbol represents an inactive asynchronous copy with errors on one or more pair.
	HyperSwap or Open HyperSwap	This symbol indicates that the HyperSwap or Open HyperSwap feature is enabled for the relationship. If a failure occurs when I/O is being written to the primary storage system, these features automatically swap the I/O to the secondary site with no user interaction and little or no application impact.
	Suspended	This symbol represents a suspended copy relationship.
	Failed over	This symbol represents a failover copy relationship.

## Session states

You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

**Attention:** Use only the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) or command-line interface CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through other interfaces such as the System Storage DS CLI. If you modify relationships through other interfaces, a loss of consistency can occur across the relationships that are managed by the session. The exceptions to this requirement are failover operations that are managed by external applications for certain session and storage system types as described in “Failover and failback operations” on page 347.

The **Refresh States** command is used to refresh the states of the role pairs that are in a session (if applicable for the session type) and refresh the state of the session if it is incorrect. Issue this command to query the states of the copy sets on the hardware. You are not required to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself. This command triggers multiple queries on the hardware, which can impact hardware performance. Do not run this command more frequently than every few minutes in each session.

The following table describes each session state.

*Table 108. Session states*

State	Session type	Description
Defined	All	The session exists but is inactive.
Flashing	All	In a Metro Mirror or Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 2.  In a Metro Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 3.



Table 108. Session states (continued)

State	Session type	Description
Prepared	All	<p>The source to target data transfer is active.</p> <p>In a Metro Mirror, Global Mirror, or Metro Global Mirror session, the data that is written to the source is transferred to the target, and all volumes are consistent and recoverable.</p> <p>In a FlashCopy session, the volumes are not yet consistent, but the flash is ready to begin.</p> <p><b>Important:</b> For sessions on the following storage systems, do not alter the relationships on the hardware that you established with Tivoli Storage Productivity Center for Replication:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V3500 (use only FlashCopy sessions for this storage system)</li> <li>• Storwize V3700</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> <p>For example, if a Metro Mirror session with one copy set is in the Prepared state, and you stop the role pair, the session is still displayed in the Prepared state.</p>
Preparing	All	<p>The volumes are initializing, synchronizing, or resynchronizing.</p> <p>In a Metro Mirror, Global Mirror, or Metro Global Mirror session, synchronization occurs after the first <b>Start</b> command is issue on a session. Resynchronization occurs when a volume was prepared and then suspended. The hardware records the changed tracks so that on the next startup, only the changed tracks are copied.</p> <p>In a FlashCopy session, the volumes are initializing. The preparing state for FlashCopy sessions applies only to SAN Volume Controller or Storwize storage systems.</p>
Recovering	All	<p>The session is in the process of recovering.</p>
Suspended	All	<p>Data copying has temporarily stopped.</p> <p><b>Important:</b> The suspended state applies only to Global Mirror, Metro Mirror, and Metro Global Mirror sessions.</p>
Suspended (Partial)	Multi-target sessions	<p>Data copying between the source site and one of the target sites has suspended.</p>

Table 108. Session states (continued)

State	Session type	Description
SuspendedH1H3 MGM	Metro Global Mirror	Data copying between site 1 and site 3 is suspended.
SuspendedH1H2 MGM	Metro Global Mirror	Data copying between site 1 and site 2 is suspended.
Suspending	All	The session is changing into a Suspended state. <b>Important:</b> The Suspending state applies only to Metro Mirror Sessions in Global Copy mode, Global Mirror, and Metro Global Mirror sessions and does not apply to SAN Volume Controller or Storwize storage systems.
Target available	All	Target volumes are available for application updates.
Terminating	FlashCopy	The session is being terminated because you issued a Terminate action under the following conditions: <ul style="list-style-type: none"> <li>You permitted the target to be Metro Mirror or Global Copy.</li> <li>You set the Require or Attempt to Preserve Mirror option.</li> </ul> The session displays as Terminating until the FlashCopy background copy is complete and no longer exists on the hardware.

## Role pair status and progress

In a session, a *role pair* is the association of two volume roles that take part in a copy relationship. Tivoli Storage Productivity Center for Replication provides detailed role pair status and progress messages for sessions.

The role pair status and progress messages are updated to provide a message that indicates what the session is doing at the time. By hovering over a progress bar, you can see specific information about the action running on the session. Some status messages might include an estimated time-to-completion for the action in hours and minutes.

Role pair status is not provided for the XIV system Snapshot sessions because role pairs are not used for these sessions.

The status messages are displayed in the Session Details and Role Pair Details page.

Table 109. Detailed status messages for Participating and Non-Participating role pairs

Supported session type	Status message
FlashCopy Metro Mirror Global Copy Global Mirror	Starting <i>role_pair_name</i> relationships on the hardware
Metro Mirror Global Mirror	Waiting for all pairs in the role pair <i>role_pair_name</i> to reach state of <i>state</i>

Table 109. Detailed status messages for Participating and Non-Participating role pairs (continued)

Supported session type	Status message
FlashCopy Metro Mirror Global Copy Global Mirror	Terminating all pairs in role pair <i>role_pair_name</i>
FlashCopy Metro Mirror Global Mirror	Recovering all pairs in role pair <i>role_pair_name</i>
Metro Mirror Global Copy Global Mirror	Suspending all pairs in role pair <i>role_pair_name</i>
FlashCopy	Background copy is running for role pair <i>role_pair_name</i>
Global Copy	Waiting for all pairs in role pair <i>role_pair_name</i> to become consistent
Global Copy	Waiting for all pairs in role pairs <i>role_pair_name</i> to complete the initial copy
FlashCopy	Waiting for all pairs in role pairs <i>role_pair_name</i> to complete FRR
Global Mirror	Waiting for all pairs in role pairs <i>role_pair_name</i> to join the Global Mirror session

### Viewing session details

You can view detailed information about a session, including role pairs, error count, whether the session is recoverable, copying progress, session type, and the timestamp.

Perform these steps to view session details:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session that you want to view.
3. From the **Session Actions** list, select **View/Modify > View Details**.

### Viewing additional details for Global Mirror sessions

Additional detail information is available for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Global Mirror sessions that do not have an Inactive status. This information includes details about the Global Mirror master, consistency groups that have been formed, and data exposure.

Complete these steps to view additional details for Global Mirror sessions:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the Global Mirror session that you want to view.
3. From the **Session Actions** list, select **View/Modify > View Details**.
4. Click the **Global Mirror Info** tab. The following information is displayed on the tab:

#### Global Mirror Master

Shows the name of the storage system acting as the Global Mirror master.

**Last Master Consistency Group Time**

Shows the time that the last consistency group was formed

**Master Time During Last Query**

Shows the time on the master storage device when the query was performed,

**Data Exposure**

Shows the average exposure to potential data loss in seconds over the query interval.

**Session ID**

Shows the Global Mirror session ID.

**Master State**

Shows the state of the master session on the hardware.

**Unsuccessful CGs since last successful CG**

Shows the number of consistency groups that have failed to form since the last successful consistency group was formed.

**CG Interval Time**

Shows the interval time between attempts to form a consistency group.

**Max Coordination Interval**

Shows the extended distance consistency maximum coordination interval.

**Max CG Drain Time**

Shows the maximum time the consistent set of data is allowed to drain at the remote site before failing consistency group formation.

**Unsuccessful CGS/Previous Query**

Shows the number of consistency groups and percentage of consistency groups that were unsuccessful since the previous query.

**Unsuccessful CGS/Total**

Shows the total number of unsuccessful consistency groups and percentage of consistency groups that have failed.

**Successful CGS/Previous Query**

Shows the number of consistency groups and percentage of consistency groups that were successful since the previous query.

**Successful CGS/Total**

Shows the total number of successful consistency groups and percentage of consistency groups that have been successful.

**Consistency Group Failure Messages**

Shows the failure messages that have occurred on the Global Mirror session that prevented the formation of a consistency group.

**Data Exposure chart**

Shows the data exposure values in seconds for the last 15 minutes or 24 hours.

**Highlight Data Exposure**

Use the following fields to define a value in seconds for which you want data exposure to tracked in the **Data Exposure** chart.

**Show Data Exposure over**

Data exposure that is over the value that is entered in this field is shown in the **Data Exposure** chart.

### Show Data Exposure under

Data exposure that is under the value that is entered in this field is shown in the **Data Exposure** chart.

## Viewing session properties

You can view and modify the description and options for a session. The options that are presented depend on the storage system type and the session type that you select.

To view or modify the properties for a session, complete the following steps:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session with the properties that you want to view.
3. From the **Session Actions** list, click **View/Modify > Properties**. The properties that are displayed depend on the session type as described in the following sections.

### TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions

The properties for IBM TotalStorage Enterprise Storage Server Model 800, IBM System Storage DS8000, and IBM System Storage DS6000 sessions depend on the session type.

The following sections describe the properties for each session type.

#### FlashCopy session properties:

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 FlashCopy sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-T1 Options** tab to set properties that apply to the role pair.

#### Session Options

Enter the properties for the session.

##### Description

Type the description for this session.

#### H1-T1 Options

Enter the properties for the role pair.

##### Incremental

Select this option to apply incremental changes to the target volume. After the initial FlashCopy operation, only data that changed on the source volume since the last FlashCopy operation was performed is copied to the target volume.

If you select this option, a persistent FlashCopy relationship is created regardless of whether you select the **Persistent** check box.

##### Persistent

Select this option to keep the FlashCopy relationship established on the hardware after all source tracks are copied to the target volume. If you do not select this option, the local replication

relationship ends after the target volume contains a complete point-in-time image of the source volume.

### **No Copy**

Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

This option is required for space-efficient volumes.

### **Allow FlashCopy target to be Metro Mirror source**

Select this option to enable the FlashCopy operation if the target volume of the FlashCopy relationship is also the source volume of a Metro Mirror relationship. If this option is not selected, the FlashCopy operation fails.

**Requirement:** This option requires that the IBM Remote Pair FlashCopy option is available for your IBM System Storage DS8000 storage system. To determine whether you can use the IBM Remote Pair FlashCopy option with your System Storage DS8000 storage system, refer to the System Storage DS8000 documentation for microcode level that you are using.

Select one of the following options to specify whether you want to maintain consistency, if possible:

#### **Do not attempt to preserve Metro Mirror consistency**

Click this option if you want the FlashCopy operation to complete without preserving consistency of the Metro Mirror relationship on the remote site. The FlashCopy operation does not occur on the remote site.

#### **Attempt to preserve Metro Mirror consistency, but allow FlashCopy even if Metro Mirror target consistency cannot be preserved**

Click this option to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship when both the source and target of the FlashCopy relationship are the source of a Metro Mirror relationship. If the consistency cannot be preserved, a full copy of the Metro Mirror relationship at the target of the FlashCopy relationship is performed. To preserve consistency, parallel FlashCopy operations are performed on both sites if possible.

#### **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved**

Click this option to prevent a full copy from being performed over the Metro Mirror link. Instead, if possible, parallel FlashCopy operations are performed on both sites. If the consistency cannot be preserved, the flash for the FlashCopy relationships fails, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship is not changed.

## **Metro Mirror Single Direction session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Single Direction sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

### **Session Options**

Enter the properties for the session.

#### **Description**

Type the description for this session.

#### **Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

#### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

### **H1-H2 Options**

Enter the properties for the role pair.

#### **Metro Mirror Suspend Policy**

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

##### **Hold I/O after Suspend**

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

##### **Release I/O after Suspend**

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect

on the applications that are writing to the source volume. This option is enabled by default.

**Attention:** This option can cause the data on the source volume to be different from the data on target volume.

### **z/OS Management**

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

#### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

#### **Enable Hardened Freeze**

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

**Requirement:** This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.



## **Metro Mirror Failover/Failback session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Failover/Failback sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

### **Session Options**

Enter the properties for the session.

#### **Description**

Type the description for this session.

#### **Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

#### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

### **H1-H2 Options**

Enter the properties for the role pair.

#### **Metro Mirror Suspend Policy**

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

##### **Hold I/O after Suspend**

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

##### **Release I/O after Suspend**

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect

on the applications that are writing to the source volume. This option is enabled by default.

**Attention:** This option can cause the data on the source volume to be different from the data on target volume.

### **z/OS Management**

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

#### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

#### **Enable Hardened Freeze**

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

**Requirement:** This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened**

**Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap are not enabled for a session.

#### **Manage H1-H2 with HyperSwap**

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Configuration Error:**

##### **Partition the system(s) out of the sysplex**

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

##### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Planned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

#### **On Unplanned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

**Requirement:** If you select this option, you must restart the system.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

## **Metro Mirror Failover/Failback with Practice session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Failover/Failback with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session or multiple role pairs in the session. Use the **H2-I2 Options** tab to set properties that apply to that role pair.

### **Session Options**

Enter the properties for the session.

#### **Description**

Type the description for this session.

#### **Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

#### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

#### **Metro Mirror Suspend Policy**

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

##### **Hold I/O after Suspend**

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

##### **Release I/O after Suspend**

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect

on the applications that are writing to the source volume. This option is enabled by default.

**Attention:** This option can cause the data on the source volume to be different from the data on target volume.

### **z/OS Management**

Select the following options to manage IBM z/OS features for the H1-I2 and H2-H1 role pairs.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connection in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

#### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

#### **Enable Hardened Freeze**

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

**Requirement:** This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

## H2-I2 Options

Enter the properties for the role pair.

### Persistent

Select this option to keep FlashCopy pairs persistent on the hardware.

This option is available only for IBM System Storage DS8000 version 4.2 or later.

## Global Mirror Single Direction and Failover/Failback session properties:

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Single Direction and Global Mirror Failover/Failback sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-J2 Options** tab to set properties that apply to the role pair.

## Session Options

Enter the properties for the session.

### Description

Type the description for this session.

### Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

### Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

## H1-J2 Options

Enter the properties for the role pair.

### Consistency group interval time (seconds)

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

### Recovery Point Objective Alerts

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

**Warning level threshold (seconds)**

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

**Severe level threshold (seconds)**

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

**Global Mirror Failover/Failback with Practice session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Failover/Failback with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, options for the H1-J2 role pair are contained in the **H1-J2 Options** tab.

**Session Options**

Enter the properties for the session.

**Description**

Type the description for this session.

**Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

**Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

**H1-J2 Options**

Enter the properties for the role pair.

### **Consistency group interval time (seconds)**

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

### **Recovery Point Objective Alerts**

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

#### **Warning level threshold (seconds)**

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

#### **Severe level threshold (seconds)**

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

### **H2-I2 Options**

Enter the properties for the role pair. The following options are available only for IBM System Storage DS8000 version 4.2 or later.

#### **Persistent**

Select this option to keep FlashCopy pairs persistent on the hardware.

#### **No Copy**

Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the H2 volume until the blocks or tracks of the I2 volume are modified.

### **I2-J2 Options**

Enter the properties for the role pair.

#### **Reflash After Recover**

Select this option if you want to create a FlashCopy replication between the I2 and J2 volumes after the recovery of a Global Mirror Failover/Failback with Practice



session. If you do not select this option, a FlashCopy replication is created only between the I2 and H2 volumes.

This option is available only for IBM System Storage DS8000 version 4.2 or later.

### **Global Mirror Either Direction with Two-Site Practice session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Either Direction with Two-Site Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, properties for the H1-J2 role pair are contained in the **H1-J2 Options** tab.

#### **Session Options**

Enter the properties for the session.

##### **Description**

Type the description for this session.

##### **Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

##### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

#### **H1-J2 Options and H2-J1 Options**

Enter the properties for the role pair.

##### **Consistency group interval time (seconds)**

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

##### **Recovery Point Objective Alerts**

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

**Warning level threshold (seconds)**

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

**Severe level threshold (seconds)**

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

**Metro Global Mirror session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Global Mirror sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, properties for the H1-H2 role pair are contained in the **H1-H2 Options** tab.

**Session Options**

Enter the properties for the session.

**Description**

Type the description for this session.

**Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

**Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

**H1-H2 Options**

Enter the properties for the role pair.

**Metro Mirror Suspend Policy**

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is

suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

#### **Hold I/O after Suspend**

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

#### **Release I/O after Suspend**

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

**Attention:** This option can cause the data on the source volume to be different from the data on target volume.

### **z/OS Management**

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

#### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Volumes in the session must be attached to the system or sysplex that you choose to enable the z/OS features. Commands for the features are issued to this system or sysplex.

#### **Enable Hardened Freeze**

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

**Requirement:** This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use

Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap are not enabled for a session.

#### **Manage H1-H2 with HyperSwap**

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Configuration Error:**

##### **Partition the system(s) out of the sysplex**

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

##### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Planned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an

error occurs during HyperSwap processing. HyperSwap is disabled.

#### **On Unplanned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

**Requirement:** If you select this option, you must restart the system.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

#### **H1-J3 Options and H2-J3 Options**

Enter the properties for the role pair.

##### **Consistency group interval time (seconds)**

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

##### **Recovery Point Objective Alerts**

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

##### **Warning level threshold (seconds)**

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

##### **Severe level threshold (seconds)**

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is

greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

### **Metro Global Mirror with Practice session properties:**

Use the View/Modify Properties notebook to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Global Mirror with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set properties that apply to specific role pairs. For example, properties for the H1-H2 role pair are contained in the **H1-H2 Options** tab.

#### **Session Options**

Enter the properties for the session.

##### **Description**

Type the description for this session.

##### **Reset Secondary Reserves**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

##### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

#### **H1-H2 Options**

Enter the properties for the role pair.

##### **Metro Mirror Suspend Policy**

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

##### **Hold I/O after Suspend**

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you

do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

### **Release I/O after Suspend**

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

**Attention:** This option can cause the data on the source volume to be different from the data on target volume.

### **z/OS Management**

Select the following options to manage IBM z/OS features for the role pair.

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

### **Enable Hardened Freeze**

Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

**Requirement:** This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing

the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap are not enabled for a session.

#### **Manage H1-H2 with HyperSwap**

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Configuration Error:**

##### **Partition the system(s) out of the sysplex**

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

##### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Planned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

#### **On Unplanned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

**Requirement:** If you select this option, you must restart the system.



### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

### **H1-J3 Options and H2-J3 Options**

Enter the properties for the role pair.

#### **Consistency group interval time (seconds)**

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

#### **Recovery Point Objective Alerts**

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

#### **Warning level threshold (seconds)**

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

#### **Severe level threshold (seconds)**

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

### **H3-I3 Options**

Enter the properties for the role pair.

#### **No Copy**

Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the H3 volume until the blocks or tracks of the I3 volume are modified.

This option is available only for IBM System Storage DS8000 version 4.2 or later.

## Metro Mirror - Metro Mirror session properties:

Use the View/Modify Properties notebook to view or modify properties for System Storage DS8000 Metro Mirror - Metro Mirror sessions.

If Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system, the properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the individual role pair option tabs to set z/OS management properties that apply to specific role pairs. For example, properties for the H1-H2 role pair are contained in the **H1-H2 Options** tab. Use the **HyperSwap Options** tab to set properties for managing HyperSwap operations for the session.

If Tivoli Storage Productivity Center for Replication is not installed on or connected to a z/OS host system, only the **Session Options** tab is displayed.

**Connecting to a z/OS host:** You can add a z/OS host connection in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

### Session Options

Enter the properties for the session.

#### Description

Type the description for this session.

#### Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

#### Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

**Attention:** This option causes the session to overwrite all data that is on the target volume.

#### Metro Mirror Suspend Policy

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume.

#### Hold I/O after Suspend

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

#### **Release I/O after Suspend**

Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the effect on the applications that are writing to the source volume. This option is enabled by default.

**Attention:** This option can cause the data on the source volume to be different from the data on target volume.

#### **System or sysplex**

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

If you select **No Association**, the options that are on the remaining tabs in the notebook are unavailable.

This option is displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system.

#### **H1-H2 Options, H1-H3 Options, and H2-H3 Options**

Enter the properties for the role pair.

#### **Manage Hx-Hx with HyperSwap**

Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2, H1-H3, or H3-H2 sequence of the session.

#### **HyperSwap Options**

Enter the HyperSwap properties for the role pairs that are in the session. If you select **Manage Hx-Hx with HyperSwap** on the **H1-H2 Options, H1-H3 Options, or H2-H3 Options** tab, these properties define how HyperSwap Manager determines the target site for the HyperSwap operation, the priority for target sites, and the process that occurs in the event of an error during the operation.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring for all of the role pairs in the session. This option is typically selected to disable HyperSwap temporarily. For example, to disable HyperSwap for planned maintenance.

#### **On Configuration Error:**

#### **Partition the system(s) out of the sysplex**

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

### **Disable HyperSwap**

Select this option to prevent a HyperSwap operation from occurring.

#### **On Planned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

#### **On Unplanned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

If you select this option, you must restart the system.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

#### **HyperSwap Site Selection**

Select one of the following options to specify how HyperSwap Manager selects the target site for a HyperSwap operation.

##### **Allow HyperSwap Manager to determine the HyperSwap site**

Select this option to enable HyperSwap to select the target site for a HyperSwap operation.

HyperSwap Manager selects the target site that keeps the most sysplex members active. If multiple target sites are equal candidates for the HyperSwap operation, HyperSwap Manager uses the **Site HyperSwap Priorities** value to select the target site. For example, if both site 1 and site 3 are equal, but the **Site HyperSwap Priorities** priority is site 3, site 2, and site 1, HyperSwap Manager selects site 3.

##### **Determine the HyperSwap site by specified priorities**

Select this option to direct HyperSwap Manager to select the target site based on the priority list that is specified by the **Site HyperSwap Priorities** option.

#### **HyperSwap Site Priorities**

Specify the priority in which HyperSwap Manager selects the target site for a HyperSwap operation.

The default value priority is site 1, site 2, and site 3. This order indicates that the first priority is site 1, the second priority is site 2, and the third priority is site 3. If site 1 is the active site, then site 1

is ignored and HyperSwap Manager selects site 2 or 3 depending on the load status and availability of the sites, with site 2 being the preferred site.

If you select **Allow HyperSwap Manager to determine the HyperSwap site**, this option is used only when multiple target sites are equal candidates for a HyperSwap operation.

If you select **Determine the HyperSwap site by specified priorities**, this option determines the order in which HyperSwap Manager selects the target site.

## **SAN Volume Controller and Storwize sessions**

The properties for IBM System Storage SAN Volume Controller and IBM Storwize sessions depend on the session type.

The following sections describe the properties for each session type with the exception of session types that include only the **Description** property.

### **FlashCopy session properties:**

Use the View/Modify Properties notebook to view or modify properties for SAN Volume Controller and Storwize FlashCopy sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-T1 Options** tab to set properties that apply to the role pair.

#### **Session Options**

Enter the properties for the session.

#### **Description**

Type the description for this session.

#### **H1-T1 Options**

Enter the properties for the role pair.

#### **Incremental**

Select this option to apply incremental changes to the target volume. After the initial FlashCopy operation, only data that changed on the source volume since the last FlashCopy operation was performed is copied to the target volume.

#### **Background copy rate (percentage)**

Type the copy rate that the storage system uses to perform the background copy of the FlashCopy role pair. You can specify a percentage in the range of 0 - 100. The default is 50.

Specify 0 if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, IBM Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

### **Metro Mirror and Global Mirror Failover/Failback with Practice session properties:**

Use the View/Modify Properties notebook to view or modify properties for SAN Volume Controller and Storwize Metro Mirror Failover/Failback with Practice and Global Mirror Failover/Failback with Practice sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H2-I2 Options** tab to set properties that apply to the role pair.

#### **Session Options**

Enter the properties for the session.

##### **Description**

Type the description for this session.

#### **H2-I2 Options**

Enter the properties for the role pair.

##### **Incremental**

Select this option to set up the relationship for recording changes to the practice volume (H2). All subsequent FlashCopy operations between the intermediate volume and the host volume copy only the data that changed since the previous FlashCopy operation.

##### **Background Copy Rate for H2-I2**

Type the copy rate that the storage system uses to perform the background copy of the FlashCopy role pair. You can specify a percentage in the range of 0 - 100. The default is 50.

Specify 0 if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, IBM Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

### **Global Mirror Failover/Failback with Change Volumes session properties:**

Use the View/Modify Properties notebook to view or modify properties for SAN Volume Controller and Storwize Global Mirror Failover/Failback with Change Volumes sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.

#### **Session Options**

Enter the properties for the session.

##### **Description**

Type the description for this session.

## H1-H2 Options

Enter the properties for the role pair.

### Enable change volumes

Select this option to enable the use of change volumes in the copy sets for the session. If this option is not selected, the session functions the same as a Global Mirror Failover/Failback session.

Change volumes are denoted as Cx, where x identifies the site. These volumes contain point-in-time images that are copied from the H1 and H2 volumes. The C1 volume stores changes from the H1 volume. These changes are sent from the C1 volume to the H2 volume, and then to the C2 volume.

Because the data that is replicated between sites contains point-in-time changes rather than all changes, a lower bandwidth link is required between the sites when change volumes are used. However, the use of change volumes can result in an increase to data exposure. Therefore, you might want to enable or disable this option depending on your network traffic or business requirements.

When you select this option, it is the equivalent of setting the **-cyclingmode** parameter to `multi` on a SAN Volume Controller or Storwize storage system. When you clear this option, it is the equivalent of setting the **-cyclingmode** parameter to `none` for these storage systems.

### Cycle period (seconds)

Specify the amount of time, in seconds, in which the change volumes will be refreshed with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. The range of possible values is 60 - 86400. The default is 300.

### Recovery Point Objective Alerts

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time, in seconds, of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 172800 seconds. The default is 0 seconds, which specifies that no alerts are generated.

#### Warning level threshold (seconds)

Specify the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Cycle period (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

#### Severe level threshold (seconds)

Specify the number of seconds that you want to set for the

severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

## **XIV system sessions**

The properties for the XIV system sessions depend on the session type.

The following sections describe the properties for each session type with the exception of session types that include only the **Description** property.

### **XIV system Global Mirror Failover/Failback session properties:**

Use the View/Modify Properties notebook to view or modify properties for the XIV system Global Mirror Failover/Failback sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** to set properties that apply to the role pair.

#### **Session Options**

Enter the properties for the session.

#### **Description**

Type the description for this session.

#### **H1-H2 Options**

Enter the properties for the role pair.

#### **Recovery point objective threshold (seconds)**

Type the number of seconds that you want to set for the recovery point objective (RPO) threshold for the role pair.

If the XIV system determines that the RPO is greater than this threshold value, an alert is generated and the session status changes to Severe. You can specify an RPO in the range of 30 - 86400 seconds. The default is 30 seconds.

#### **Synchronization schedule (HH:MM:SS)**

Select an interval to create a synchronization schedule. The schedule is used to form consistency groups. The XIV system attempts to form consistent points of data by taking automatic snapshots of the volumes in the session at this interval. The default is **Minimum Interval**, which is 20 seconds.

If you select **Never**, synchronization is not scheduled and the XIV system does not create consistency groups. When the XIV system determines that the RPO threshold is exceeded, the session state becomes Severe.

## **Basic HyperSwap session properties**

Use the View/Modify Properties notebook to view or modify properties for Basic HyperSwap sessions.

The properties are displayed on different tabs in the notebook. Use the **Session Options** tab to set properties that apply to the entire session. Use the **H1-H2 Options** tab to set properties that apply to the role pair.



## Session Options

Enter the properties for the session.

### Description

Type the description for this session.

### Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

## H1-H2 Options

Enter the properties for the role pair.

### z/OS Management

The **z/OS Management** options are displayed only if Tivoli Storage Productivity Center for Replication is installed on or connected to a z/OS host system. You can add a z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

When you are connected to a z/OS host system, you can manage z/OS features for volumes that are attached to the host system. In addition, you can manage these features for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex.

### System or sysplex

Select the z/OS system or sysplex that you want to associate with the session. Commands for z/OS features are issued to this system or sysplex. Volumes in the session must be attached to the system or sysplex to enable the features.

### Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

### On Configuration Error:

#### Partition the system(s) out of the sysplex

Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

#### Disable HyperSwap

Select this option to prevent a HyperSwap operation from occurring.

### On Planned HyperSwap Error:

#### Partition out the failing system(s) and continue swap processing on the remaining system(s)

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

#### Disable HyperSwap after attempting backout

Select this option to enable IOS to back out the

HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

#### **On Unplanned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

**Requirement:** If you select this option, you must restart the system.

##### **Disable HyperSwap after attempting backout**

Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

## **Viewing copy sets in a session**

You can view the copy sets that are in a session. You can also export the copy sets to a comma-separated value (CSV) file for backup purposes.

Follow these steps to view or export the copy sets for a session:

1. In the menu bar, click **Sessions**.
2. On the Sessions page, select the session that contains the copy sets.
3. From the **Session Actions** list, select **View/Modify > View Copy Sets**.

## **Viewing storage system details**

You can view detailed information about storage system, including the name, location, type, vendor, and the status of all connections to the storage system.

Perform these steps to view storage system details:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click the ID for the storage system in **Storage System** column.

## **Viewing storage connection details**

You can view storage connection details and a list of all storage systems that are located behind the connection.

Perform these steps to view storage connection details:

1. In the menu bar, click **Storage > Storage Systems**.
2. On the Storage Systems page, click the **Connections** tab.
3. Click the ID for the storage system connection in the **Storage Connection** column.

## **Viewing volume details**

You can view information about volumes such as the name of the volume, the capacity of the volume, and the type of volume.

1. In the menu bar, click **Storage > Volumes**.
2. On the Volumes page, select a storage system.
3. Depending on the type of storage system, complete one of the following actions:
  - Select **All IO Groups** or a specific I/O group.
  - Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
  - Select **All Pools** or a specific pool.
4. Click **Perform Query**. Information about the volumes is displayed in a table.

## Viewing logical paths

You can view all logical paths that are defined on an TotalStorage Enterprise Storage Server, System Storage DS8000, or System Storage DS6000 storage system.

Complete one of these procedures to view logical paths:

- From the ESS/DS Paths page of Tivoli Storage Productivity Center for Replication:
  1. In menu bar, click **Paths**.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems page:
  1. In menu bar, click **Storage > Storage Systems**.
  2. Select the storage system for which you want to view logical paths.
  3. From the **Select Action** list, select **View Paths**. The paths page is displayed with a list of defined logical paths.

## Viewing host connection details

You can view and modify the connection information for a host system. You can modify the connection information only if the host system in the Disconnected state.

Follow these steps to view or modify connection information for a host system:

1. In the menu bar, click **Storage > Host Connections**.
2. On the Host Connections page, click the link for the host system in the **Host System** column.

## Viewing console messages

Tivoli Storage Productivity Center for Replication provides a detailed log of user and system activity.

To view this log, click **Console** in the menu bar. The Console window contains log entries that specify the activity that has occurred and a link to the correlating message for that activity.

---

## Security

The IBM Tivoli Storage Productivity Center for Replication authentication process uses a configured user registry from either the operating system or Lightweight Directory Access Protocol (LDAP) server. To perform a specific action and manage specific sessions in the IBM Tivoli Storage Productivity Center for Replication GUI or CLI, the user must also have the appropriate authorization. Authorization is granted by assigning a specific role to the user account or user group.

## Users and groups

To log on to Tivoli Storage Productivity Center for Replication, user ID and password is required.

The information in this topic is specific to Tivoli Storage Productivity Center for Replication on open systems. For information about defining users and groups for Tivoli Storage Productivity Center for Replication for System z, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

Tivoli Storage Productivity Center for Replication does not maintain a directory of user names and passwords. Instead, the application uses one of the following repositories for user authentication:

- An operating system repository
- A Lightweight Directory Access Protocol (LDAP) repository

Tivoli Storage Productivity Center uses the federated repositories feature in WebSphere Application Server and an operating system repository is created by default during installation of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.

If you choose to use LDAP authentication, you must add the LDAP repository to the federated framework after the installation of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication is complete.

For more information about the operating system and LDAP repositories, see the topic about changing the user authentication configuration in the Tivoli Storage Productivity Center documentation.

### **Adding Tivoli Storage Productivity Center users and groups to Tivoli Storage Productivity Center for Replication**

By default, the common user that was defined during the installation of Tivoli Storage Productivity Center and the user `tpcFileRegistryUser` are granted access to Tivoli Storage Productivity Center for Replication and assigned to the Administrator role.

The user `tpcFileRegistryUser` is used only for recovery purposes. For example, if you accidentally delete the repository that you are using for authentication, you can access Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication by using the user `tpcFileRegistryUser`.

The password for the user `tpcFileRegistryUser` is the same as the password that was entered for the common user during the installation of Tivoli Storage Productivity Center.

To ensure smooth integration of Tivoli Storage Productivity Center with Tivoli Storage Productivity Center for Replication, complete the following actions:

- Add all Tivoli Storage Productivity Center users and groups (other than those that are assigned by default) to Tivoli Storage Productivity Center for Replication. For example, if you added a `TPCSuperuser` group to Tivoli Storage Productivity Center for LDAP authentication, add that group to Tivoli Storage Productivity Center for Replication as well.
- Use the same user or group to log on to both applications.

## User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

Restricting access to sessions prevents unwarranted administrative access. This is especially useful in an open environment, where there can be many storage administrators who are responsible for their servers, applications, databases, file systems, and so on.

Tivoli Storage Productivity Center for Replication provides a set of predefined user roles: Monitor, Operator, and Administrator.

### Monitor

Monitors can view the health and status in the Tivoli Storage Productivity Center for Replication GUI and CLI; however, they cannot modify or perform any commands or actions.

Monitors can view the following information:

- All storage systems and storage system details
- All connections and connection details
- All sessions and session details
- All path information
- Management server status and details

### Operator

Operators can manage sessions to which they have been assigned, including:

- Adding or removing a session. The user ID that created the session is automatically granted access to manage that session.
- Performing actions on an assigned session, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding copy sets to a session. The session operator can add volumes to a copy set only when the volume is not protected and not in another session.
- Removing copy sets from a session.
- Adding Peer To Peer Remote Copy (PPRC) paths, and removing paths with no hardware relationships. PPRC paths are a common resource used in Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage system relationship that is established between two common logical subsystems (LSSs).

**Note:** The session operator cannot issue a force removal of a path.

**Note:** A path can also be auto-generated when starting a session.

- Monitoring health and status, including viewing the following information:
  - All storage systems and storage system details
  - All connections and connection details
  - All sessions and session details
  - All path information
  - Management server status and details

**Note:** Operators can make changes only to the volumes that they own. They are not able to make changes to volumes being managed by other users.

## Administrator

Administrators have unrestricted access. They can manage all sessions and perform all actions associated with Tivoli Storage Productivity Center for Replication, including:

- Granting permissions to users and groups of users.
- Adding or removing a session. The user ID that created the session is automatically granted access manage that session.
- Performing actions on all sessions, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding and removing copy sets from a session. The administrator can add volumes to a copy set only when the volume is not protected and not in another session.
- Protecting volumes and removing volume protection.
- Adding or removing storage system connections.
- Modifying connection properties.
- Assigning or changing storage system locations.
- Adding PPRC paths and removing paths with no hardware relationships. PPRC paths are a common resource used in Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

**Note:** A path can also be auto-generated when starting a session.

- Managing management servers. The standby management server is a common resource that is available to multiple sessions.
- Packaging program error (PE) log files.
- Monitoring health and status, including viewing the following information:
  - All storage systems and storage system details
  - All connections and connection details
  - All sessions and session details
  - All path information
  - Management server status and details

**Note:** Administrators cannot revoke their own administrative access rights.

## Granting access privileges for a user

You can assign user roles to a Tivoli Storage Productivity Center for Replication user to grant access privileges to individual sessions and tasks.

Perform the following steps to authorize a user:

1. Log on to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings > Administration**.
3. On the Administration page, click **Add Access**. The Add Access wizard is displayed.
4. Type the name of the user to whom you want to give access, and click **Next**. The Select Users and Groups page is displayed.

**Tip:** You can enter a partial name and use the \* wildcard character to represent zero or more characters.

5. Select one or more names from the list of found users.
6. Select the role to associate with this user.
7. If you selected the Operator role, select one or more session that this user can manage, and click **Next**.
8. Click **Next** to confirm this action.
9. Click **Finish**.

## Viewing access privileges for a user

You can view a list of all Tivoli Storage Productivity Center for Replication user and their assigned roles. You can also view the assigned sessions for each user.

Complete the following steps to view access privileges for a user:

1. Log in to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings > Administration**. The Administration page is displayed with a list of Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user whose access privileges you want to view.
4. From the **Select Action** list, select **View/Modify Access**. The View/Modify Access page is displayed. This page shows the role assigned to the user and lists the sessions that the user can manage.
5. Click **Cancel** to close the page.

## Modifying access privileges for a user

You can change the user role and assigned sessions for a Tivoli Storage Productivity Center for Replication user.

**Prerequisite:** You must have Administrator privileges to perform this action.

Complete the following steps to modify the access privileges for a user:

1. Log in to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings > Administration**. The Administration page is displayed with a list of Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user whose access privileges you want to modify.
4. From the **Select Action** list, select **View/Modify Access**. The View/Modify Access page is displayed. This page shows the role assigned to the user and lists the sessions that the user can manage.
5. Select the role to associate with the user.

6. If you selected the Operator role, select one or more sessions that the user can manage and click **Next**.
7. Click **OK**.

## Removing access privileges for a user

You can remove access privileges for a Tivoli Storage Productivity Center for Replication user. When you remove access, the user ID cannot access the Tivoli Storage Productivity Center for Replication GUI or run commands from the command line.

**Prerequisite:** You must have Administrator privileges to perform this action.

Perform the following steps to remove user access:

1. Log in to Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the menu bar, click **Settings > Administration**. The Administration page is displayed with a list of Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user for which you want to remove access.
4. From the **Select Action** list, select **Remove Access**.



---

## Chapter 7. Reporting

Tivoli Storage Productivity Center provides multiple user interfaces for viewing reports about the storage infrastructure in an enterprise environment.

### IBM Cognos reporting interface

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in the Cognos software. The Cognos software is included in Tivoli Common Reporting.

**Restriction:** Cognos Workspace is not supported by Tivoli Storage Productivity Center reports in this release. If you use Tivoli Storage Productivity Center reports to create workspaces, those workspaces might be incompatible with future releases of the product.

### vSphere Web Client extension interface

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices. You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems. To view Tivoli Storage Productivity Center storage information in block and file storage reports, you must register Tivoli Storage Productivity Center as a VASA provider. To view information about reporting on your virtual environment and storage devices, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_c\\_vm\\_reporting.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_c_vm_reporting.html).

### Stand-alone GUI

View detailed information about the storage resources in your environment in the stand-alone GUI. These reports are organized into different types and categories and provide both summary and detailed information, depending on your needs. Many reports are also accessible through the topology viewer, which provides a visual representation of storage topology. To view information about reporting in the stand-alone GUI, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_c\\_reporting.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_c_reporting.html).

---

## Getting started

You can view predefined reports and create custom reports about Tivoli Storage Productivity Center in the IBM Cognos reporting software.

You access reports from the Tivoli Storage Productivity Center web-based GUI, and work with the reports in the Cognos software. Tivoli Common Reporting provides the Cognos software for Tivoli Storage Productivity Center reports.

## Predefined reports

You can use IBM Cognos reporting software to view predefined reports about Tivoli Storage Productivity Center. Predefined reports provide information about the capacity and performance of your resources.

### Predefined reports

You can view the following types of predefined reports:

#### Reports about the capacity and relationships of your resources

Reports about the capacity and relationships of servers, hypervisors, network-attached storage systems, storage virtualizers, storage systems, storage resource groups, groups, fabrics, and switches. Use these reports to analyze the capacity and relationships of resources. For example, to show the current space statistics for a pool on a storage system, run the Pools Capacity report. To show the pools that are related to a storage system, run the Storage Resource Relationships Summary report.

#### Reports about performance

Reports about the performance of your monitored resources. You can view reports about storage systems, components of storage systems, switches, and switch ports. Use these reports to analyze the performance of resources. For example, to find out which are the most active volumes on a storage system over a specified period, run the Most Active Volumes report. To view details about the performance of a volume over a specified period, run a report such as the Performance of One Volume report.

#### Reports about historical capacity

Reports that include the historical information about the capacity of your storage systems, pools, and volumes. Use these reports to determine whether more storage resources are required. For example, to show how space is used over a specified period on volumes on storage systems, run the Volumes Historical Capacity report. You can also use these reports to analyze future storage requirements for your environment.

#### Reports about storage tiering

Reports that include information about the workload activity of these resources: managed disk (MDisk) groups, virtual disks (VDisks) in MDisk groups, VDisks in storage resource groups, and VDisks in storage virtualizers. Use these reports to determine the best storage-tier configuration for your environment.

Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics.

**Tip:** When you view a predefined report, click **Help** to find out more about the report.

#### Related reference:

“Predefined reports listed by package” on page 485

The data that the predefined reports use is organized into four packages: Capacity and Relationships, Historical Capacity, Performance, and Storage Tiering. Each report uses data from only one package.

## Predefined reports listed by package

The data that the predefined reports use is organized into four packages: Capacity and Relationships, Historical Capacity, Performance, and Storage Tiering. Each report uses data from only one package.

**Tip:** When you view a predefined report, click **Help** to find out more about the report.

### Capacity and Relationships package

You can view the following predefined reports about capacity and relationships:

#### Availability of Switch Ports

Shows a chart of the 20 switches with the most ports and connected ports, and a table with the properties of the switches.

#### Capacity of One Group

Shows charts and tables of the capacity, available space, and used space statistics for each type of resource in a group. The report shows separate charts and tables for block-level and file-level resources.

#### Disks Capacity

Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

#### File System to Volume Relationships

Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

#### File Systems Capacity

Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

#### Groups Capacity

Shows a chart of the 20 groups with the most capacity, available space, or used space, and a table with details of all groups. Use the report to review the space allocation in your groups.

#### Hypervisor Data Stores Capacity

Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

#### Hypervisor Disks Capacity

Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

#### Hypervisors Capacity

Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

#### Managed Disks Capacity

Shows a chart of the 20 managed disks with the most available space, and

a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

**Pools Capacity**

Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

**Server Disks Capacity**

Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

**Servers Capacity**

Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

**Storage Resource Relationships Summary**

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

**Storage Resource Relationships Summary (Configurable)**

Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

**Storage Systems Capacity**

Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

**Volumes Capacity**

Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

## **Historical Capacity package**

You can view the following predefined reports about historical capacity:

**Pools Historical Capacity**

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

**Storage Systems Historical Capacity**

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

**Volumes Historical Capacity**

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

## **Performance package**

You can view the following predefined reports about performance:

**Compare Performance of Multiple Controllers and Modules**

Shows up to four performance metrics for multiple controllers or modules

over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

**Compare Performance of Multiple Disks**

Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

**Compare Performance of Multiple IO Groups**

Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

**Compare Performance of Multiple Managed Disks**

Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

**Compare Performance of Multiple Nodes**

Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

**Compare Performance of Multiple Pools**

Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

**Compare Performance of Multiple Ports**

Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

**Compare Performance of Multiple Storage Systems**

Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

**Compare Performance of Multiple RAID Arrays**

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

**Compare Performance of Multiple Switch Ports**

Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

**Compare Performance of Multiple Switches**

Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

**Compare Performance of Multiple Volumes**

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

**Compare Performance of One Pool over Time Ranges**

Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

**Compare Performance of One Storage System over Time Ranges**

Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

**Compare Performance of One Switch over Time Ranges**

Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

**Compare Performance of One Switch Port over Time Ranges**

Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

**Most Active Controllers or Modules**

Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

**Most Active Disks**

Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

**Most Active Host Connections**

Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

**Most Active Hypervisors**

Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

**Most Active IO Groups**

Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

**Most Active Managed Disks**

Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

**Most Active Nodes**

Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

**Most Active Pools**

Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

**Most Active Ports**

Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

**Most Active RAID Arrays**

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

**Most Active Servers**

Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

**Most Active Storage Systems**

Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

**Most Active Switch Ports**

Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

**Most Active Volumes**

Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

**Performance Data Export**

Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM.

**Performance of One Controller or Module**

Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

**Performance of One Disk**

Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

**Performance of One IO Group**

Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

**Performance of One Managed Disk**

Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

**Performance of One Node**

Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

**Performance of One Pool**

Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

**Performance of One Port**

Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

**Performance of One RAID Array**

Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

**Performance of One Storage System**

Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

**Performance of One Switch**

Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

**Performance of One Switch Port**

Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

**Performance of One Volume**

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

**Performance of Volumes by Host Connection**

Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

**Performance of Volumes by Hypervisor**

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

**Performance of Volumes by Server**

Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

**Summarized Performance of Volumes by Host Connection**

Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

**Summarized Performance of Volumes by Hypervisor**

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

**Summarized Performance of Volumes by Server**

Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

**Storage Tiering package**

You can view the following predefined reports about storage tiering:

**MDisk Group - VDisk Workload Activity Details**

Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.



### **MDisk Group Details**

Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

### **MDisk Groups - VDisk Workload Activity**

Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDIs in the MDisk group.

### **MDisk Groups - Workload Activity**

Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

### **Storage Resource Group - VDisk Workload Activity**

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDIs in the group.

### **Storage Resource Group - VDisk Workload Activity Details**

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDIs in the group.

### **Storage Virtualizer - VDisk Workload Activity**

Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

### **VDisk Details**

Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

## **Custom reports**

You can use IBM Cognos reporting software to create custom reports about Tivoli Storage Productivity Center. Use custom reports to analyze the capacity and performance of your resources.

### **Custom reports**

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To create a report, browse the list of monitored resources in your storage environment, and drag items that are associated with the resources into a report. For example, drag items that show capacity data for storage systems into a report.

If you create custom reports, you can also do the following tasks:

- Copy a predefined report and modify the copy to create a custom report.
- Create custom charts. You can create charts with data that is in the combinations that you want to see. Use these custom charts to troubleshoot and to report on the status of resources.
- Create reports with multiple charts. You can create multiple charts with different scales, and display those charts side-by-side. You can use these charts to

compare performance metrics. For example, you can review read I/O rates and the number of bytes transferred in the same report.

- Create a custom folder structure to organize your custom reports.
- Run a report from a URL. You can include options for the report in the URL. For options that are not included in the URL, you can specify default values, or prompt the user to specify a value.
- Share a custom report with your organization, or keep it for only your own use. You can also distribute different sections of a report to different teams.

## Advanced formatting for custom reports

You can create reports with basic formatting, or design more complex reports with advanced formatting.

### Access related reports

You can configure custom reports so that you can click a resource in the report output to access other related reports. If you identify a problem in a report, you can open another report from within that report that might provide specific information about the problem.

For example, the Most Active Servers predefined report shows data for servers. Click the name of a server in the report to view the Summarized Performance of Volumes by Server report for the server. This report shows the aggregated data for all of the volumes on the server. Click the name of a volume in the report to view the Performance of Volumes by Server report for the volume.

### Use GUI controls in reports

You can configure custom reports to include GUI controls to enable users to select options for the reports. You can specify default values for these options. In addition, you can create copies of reports with different default values.

The GUI controls that you can add to your reports enable users to select dates, type values, and select from lists. You can also create GUI controls whose values depend on what is selected in another GUI control. You can include GUI controls on the report output in HTML format so that users can change options and run the report again. For example, you can include a drop-down list in the HTML report output so that users can change the sort order of the data.



### Schedule reports

You can schedule when to run the report. You can specify the following formats for the report: CSV, Excel, HTML, PDF, and XML. You can specify to email the report output and to save the report output to a location on your file system or on a web server. For example, you might require a monthly report that shows the usage of space on storage systems. Management staff might use this report to analyze changes in your storage system requirements. You can create this report and schedule it to run every month. You can specify that the report output is emailed to management staff, or saved to a location from which management staff can access it.

## Viewing predefined reports

Use the Tivoli Storage Productivity Center web-based GUI to access predefined reports.

If you are using Internet Explorer 11, you must set the browser to compatibility mode.

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Select **Cognos > View predefined reports**.
3. Navigate to the report that you want to see, and then click the name of the report. To view a description of a report or folder, point to the report icon  that is next to the report name. To view descriptions of all reports and folders, click the **Details View** icon .

## Creating custom reports


Use the Tivoli Storage Productivity Center web-based GUI to access custom reports. You can create basic reports with limited formatting, or design custom reports to add advanced formatting features to reports.

### Restriction:

You cannot create custom reports from the following packages:

- Historical Capacity
- Storage Tiering

**Note:** If you are using Internet Explorer 11, you must set the browser to compatibility mode.

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the type of report that you want to see, open the appropriate report tool:
  - To create a basic report with limited formatting about the capacity and relationships of resources, select **Cognos > Create capacity reports**.
  - To create a basic report with limited formatting about the performance of resources, select **Cognos > Create performance reports**.
3. Optional: To create a report with advanced formatting features, complete the following steps:
  - a. Click the **Return** icon .
  - b. Select **Launch > Report Studio**.

### Related tasks:

“Creating custom performance reports” on page 799

You can create reports that show performance metrics. Performance reports can contain detailed information about the performance of monitored resources and some of the properties of those resources. You access custom performance reports from the Tivoli Storage Productivity Center web-based GUI, and you create the reports in the IBM Cognos reporting software.

“Creating custom capacity and relationship reports” on page 927

You can create custom reports that show capacity and relationship information for resources. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources. You access custom capacity and relationship reports from the Tivoli Storage Productivity Center web-based GUI, and you create the reports in the IBM Cognos reporting software.

“Adding filters to reports” on page 929


You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your

capacity and relationship reports when you create the reports.

## Viewing and creating reports in the Cognos software

You access reports from the Tivoli Storage Productivity Center web-based GUI, and work with the reports in the IBM Cognos reporting software.

When you view and create reports, you can navigate from predefined reports to custom reports and from custom reports to predefined reports within the Cognos software. You do not have to return to the web-based GUI.


1. In the Cognos software, click the **Home** icon  to display the **Public Folders** tab.
2. Depending on the type of report that you want to work with, complete one of the following actions on the **Public Folders** tab:
  - To view predefined reports, click **IBM Tivoli Productivity Center Predefined Reports**.
  - To create a basic custom report with limited formatting, complete the following steps:
    - a. Click **IBM Tivoli Productivity Center Packages**.
    - b. Click the package from which you want to create a report.
    - c. Select **Launch > Query Studio**.
  - To create a report and apply advanced formatting features to the report, complete the following steps:
    - a. Click **IBM Tivoli Productivity Center Packages**.
    - b. Click the package from which you want to create a report.
    - c. Select **Launch > Report Studio**.

## Searching the predefined reports for a property or performance metric

You can search the contents of the predefined reports. You can check which reports contain a property or performance metric. For example, you can list the reports that contain the Storage Pool Name property or the Port Receive I/O Rate performance metric.

Before you can search the content of predefined reports, a search index must be created for the reports.

**Tip:** If there is no search index, you can search only the titles and descriptions of the reports, not the contents. To search the titles and descriptions of reports, select **Name field**, **Description field**, or **Name or description field** from the search options list.

1. In the search field beside the **Home** icon  in the IBM Cognos reporting software, type the name of the property or metric in quotation marks. For example, type "Storage Pool Name" or "Port Receive I/O Rate".
2. Select **Full text and all fields** from the search options list.
3. Click the search icon. The list of reports that contain the property or metric is displayed in the search results.
4. Optional: Alternatively, you can search the online help for properties or metrics. To search the online help, complete the following steps:


- a. In the Cognos software, click any predefined report.
- b. Click the **Help** link at the upper right of the page.
- c. In the **Search** field, type the name of the property or metric in quotation marks, and then click **Go**. The list of help topics that contains the property or metric is displayed in the search results.

**Related tasks:**

“Creating a search index of reports, report views, and folders” on page 507

You can create a search index of the reports, report views, and folders in the IBM Cognos reporting software.

**Related information:**

 Searching for entries using full text and all fields

For more information about searching for items in Tivoli Common Reporting, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cc.10.2.0.doc/t\\_search\\_indexed\\_search.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cc.10.2.0.doc/t_search_indexed_search.html).

---

## Scenarios

Use these scenarios to learn how to use Tivoli Storage Productivity Center reports to analyze the most active volumes on hypervisors and to monitor the performance of a critical application. You can learn how to investigate a degradation in the performance of a storage pool. You can also learn how to identify the relationships between a storage system and other resources.

**Related information:**

 IBM Tivoli Storage Productivity Center documentation

For information about the concepts behind the tasks that are associated with these scenarios, see the Tivoli Storage Productivity Center documentation at <http://www.ibm.com/support/knowledgecenter/SSNE44/>.

## Analyzing the volumes on the most active hypervisors

In this scenario, you use the predefined reports for hypervisors to analyze the volumes on hypervisors. You can view the most active volumes. Then, you can view summary information for a hypervisor and details for individual volumes. You can also use this scenario to analyze the volumes on servers, agentless servers, and cluster resource groups.

You have 10-20 volumes on the hypervisors in your storage environment. You receive reports that the performance of some volumes is slow. You want to analyze the performance of these volumes as a group, and the performance of individual volumes.

1. Run the Most Active Hypervisors report. Use the chart in the report output to see the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load. You see that the volumes on one hypervisor are considerably more active than the volumes on other hypervisors. You want to investigate that hypervisor.
2. To analyze the activity on this hypervisor, click the name of the hypervisor in the report table. The Summarized Performance of Volumes by Hypervisor report runs for the hypervisor. Use the report to see summarized performance metrics for the volumes on the hypervisor. You can also use the report to see the aggregate load and the average response time of the volumes on the hypervisor. In the charts and the table on the report, you see that the I/O response times for the hypervisor are slower than normal.

3. To investigate the individual volumes on the hypervisor, click the name of the hypervisor in the report table. The Performance of Volumes by Hypervisor report runs for the hypervisor. Use the report to see the performance metrics for individual volumes on the hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment. In the charts and the table on the report, you see that the I/O response times for one volume are slow.

You identified the volume with the slow I/O response times. You can now copy data to the other volumes to more evenly distribute the load of the volumes.

**Related reference:**

“Most Active Hypervisors report” on page 654

Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

“Summarized Performance of Volumes by Hypervisor report” on page 655

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.



“Performance of Volumes by Hypervisor report” on page 657

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

## Monitoring daily the performance of volumes

In this scenario, you create a report view from a predefined report to run a daily report about the performance of volumes. You can also use this scenario to monitor any other resources and performance metrics.

You have applications that are critical to your business. The applications access volumes in your storage environment. You want to monitor daily the performance of the volumes to confirm that they are functioning correctly. To monitor the volumes, you must create a report view to run a daily report about the performance of volumes.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.
2. Click **Storage Systems**, and then click **Volumes**.
3. Click the **Report View** icon  for the Compare Performance of Multiple Volumes report.
4. Type a name for the report view, then click **Finish**.
5. Click the **Schedule** icon  for the report view.
6. Check all seven days of the week on the **By Week** tab.
7. In the **Options** area, check **Override the default values**, and select **HTML** and **PDF** formats.
8. In the **Prompt values** area, check **Override the default values**, and then click **Set**.
9. Specify the storage systems and the volumes that you want to monitor, and then click **Next**.
10. Select the following performance metrics:
  - Total Overall I/O Rate (ops/s)
  - Overall Response Time (ms/op)

- Total Data Rate (MiB/s)
  - Total Overall Cache Hit Percentage
11. Specify an interval and a reporting period for the report view, and then click **Finish**.
  12. Click **OK** to save the schedule, format, and prompt value settings.

View the output of this report every day to confirm that these volumes are functioning correctly. Click the report view to view the latest output for the report. If the report output is not displayed, make sure that the **Default action** for report view in the properties of the report view is set to **View most recent report**.

**Related reference:**

“Compare Performance of Multiple Volumes report” on page 760

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

## Investigating a degradation in the performance of a storage pool

You can use the Compare Performance of One Pool over Time Ranges report to analyze a change in the performance of a storage pool. You can also use this scenario to compare the performance metrics of storage systems, switches, and switch ports over two periods of time.

In the previous three days, the performance of a storage pool has significantly degraded. You want to investigate the reasons for the degradation. To do this, use the Compare Performance of One Pool over Time Ranges report to compare the recent poor performance of the pool with a period when the performance was normal.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.
2. Click **Storage Systems**, and then click **Pools**.
3. Click **Compare Performance of One Pool over Time Ranges**.
4. Specify the storage system and the pool whose performance has degraded.
5. Specify the performance metric that you want to analyze, for example **Overall Response Time (ms/op)**.
6. Specify an interval.
7. Specify two custom date ranges for the two reporting periods. Specify the previous three days as one reporting period. For the other reporting period, specify a period when the performance of the pool was normal.
8. Click **Finish**.

To investigate other possible reasons for the degradation, run the report again, and select another metric.

**Related reference:**

“Compare Performance of One Pool over Time Ranges report” on page 736

Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

## Identifying the relationships between a storage system and other resources

In this scenario, you use a predefined relationships report to identify the relationships between a storage system and other servers or hypervisors. You can also use this scenario to identify the relationships between storage systems and NAS, pools, and volumes. Use this scenario to identify the relationships between storage virtualizers and servers, hypervisors, NAS, pools, and volumes.

You want to take a storage system offline and migrate the volumes to another storage system. To prepare for the migration, you first identify which servers or hypervisors are affected when the storage system goes offline.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.
2. Click **Storage Resource Relationships Summary**.
3. From the **Storage systems** list, select the storage system that you plan to take offline.
4. Click **Finish**.

View the output of the report to see which servers or hypervisors are affected when the storage system goes offline. You can now migrate the volumes.

### Related reference:

“Storage Resource Relationships Summary report” on page 628

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

---

## Configuring



You can configure Tivoli Storage Productivity Center reports according to your requirements.

### Setting options for reports

When you schedule or run a report, you can specify the output format and how the report is delivered.

You can select output formats for the report. For example, you can select PDF, HTML, and other formats. If you choose PDF as the format for your report, you can set options such as the orientation and paper size for reports.

You can also specify to print the report, to send the report by email, and other delivery options.

1. Navigate to a report in the IBM Cognos reporting software.
2. Depending on whether you want to schedule a report or run a report, do one of the following steps:
  - To schedule a report, and select format and delivery options, click the **Schedule** icon .
  - To run a report, and select format and delivery options, click the **Run with options** icon . To set advanced options, click **advanced options**, and then select **Run in the background**.

### Related information:



### Schedule an Entry

For more information about how to schedule a report, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.svg.ba.cognos.ug\\_cra.10.2.0.doc/t\\_createschedules.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.svg.ba.cognos.ug_cra.10.2.0.doc/t_createschedules.html).

### Report Formats

For more information about report formats, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.svg.ba.cognos.ug\\_cc.10.2.0.doc/c\\_specifyreportformat.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.svg.ba.cognos.ug_cc.10.2.0.doc/c_specifyreportformat.html).

### Distributing Reports

For more information about delivery options for reports, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.svg.ba.cognos.pp\\_usr\\_gd.10.2.0.doc/c\\_distributereports.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.svg.ba.cognos.pp_usr_gd.10.2.0.doc/c_distributereports.html).

### Set Advanced Report Options for the Current Run

For more information about advanced options for reports, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.svg.ba.cognos.ug\\_cc.10.2.0.doc/t\\_setadvancedrunoptions.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.svg.ba.cognos.ug_cc.10.2.0.doc/t_setadvancedrunoptions.html).



## Scheduling a report to run with predefined options

You can schedule a report to run on different systems and with different options that you define. To monitor the performance of different sets of volumes, you can schedule a Compare Performance of Multiple Volumes report for each of the sets.

You can specify schedules and output formats for the reports, and different ways to share the reports.

To schedule a report to run with parameters that you specify, you first create a report view from the report. Then, you schedule the report and set the parameters and options in the report view.

If the report from which a report view was created is updated, the report view is also updated.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.
2. Click the **Report View** icon  for the report.
3. Type a name for the report view, then click **Finish**.
4. Click the **Schedule** icon  for the report view.
5. In the schedule area, specify when and with what frequency you want the report to run.
6. In the **Options** area, check **Override the default values**, and select the formats for the report.
7. In the **Prompt values** area, check **Override the default values**, and then click **Set**.
8. Specify parameters for the report. For example, you can specify the resources, the interval, and the reporting period for the report.
9. Click **OK** to save the schedule, format, and prompt value settings.

After the report runs for the first time, click the report view to view the latest output for the report. If the report output is not displayed, make sure that the **Default action** for report view in the properties of the report view is set to **View most recent report**.

### Related tasks:

“Monitoring daily the performance of volumes” on page 496

In this scenario, you create a report view from a predefined report to run a daily report about the performance of volumes. You can also use this scenario to monitor any other resources and performance metrics.

### Related information:

 [Create a Report View](#)




For more information about report views, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cc.10.2.0.doc/t\\_createareportview.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cc.10.2.0.doc/t_createareportview.html).

## Creating a copy of a predefined report


Create a copy of a predefined report if you want to modify report. You might want to modify the queries that the report uses, or to add or remove a column. Changes to the predefined report do not affect copies of the report.

When you create a copy of a report, the links to related reports in the report output do not work correctly.

If the predefined report is changed during an upgrade, the copy of the report is not affected. If you modify a predefined report instead of creating a copy and then you upgrade your reports software, your modifications might be overwritten.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.
2. Navigate to the report that to want to copy, and then click the check box that is next to the title of the report.
3. Click the **Copy** icon .
4. Optional: Navigate to a folder where you want to store the copy of the report. For example, **My Folders**.
5. Click the **Paste** icon .
6. Optional: To change the name of the report, click the **Properties** icon , type the name in the **Name** field, and then click **OK**.

### Related information:

 [Copy an Entry](#)

For more information about copies of reports, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cc.10.2.0.doc/t\\_copyanentry.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cc.10.2.0.doc/t_copyanentry.html).

## Starting the Cognos Configuration GUI

You can use the IBM Cognos Configuration GUI to configure Cognos reporting software components. For example, you can configure an archive location for the output of reports, and configure a connection between the Cognos software and an SMTP server.

To start Cognos Configuration GUI, follow the instructions for your operating system:

### Windows

Click **Start > Tivoli Common Reporting 3.1.0.2 > Cognos Configuration**.

Alternatively, you can run the `tcrcogconfig.bat` file in the `JazzSM_install_directory\reporting\cognos\bin64` directory.

## AIX or Linux

Log on to a graphical session, and then run the `tcrcogconfig.sh` script in the `JazzSM_install_directory/reporting/cognos/bin64` directory.

## Sending reports by email

To send reports as attachments in email messages, configure a connection between the IBM Cognos reporting software and an SMTP server.

If you use Red Hat Enterprise Linux 5, you must set the `LD_PRELOAD` environment variable to open the IBM Cognos Configuration GUI. To set the environment variable, issue the following command:

```
setenv LD_PRELOAD /usr/lib64/libfreetype.so
```

You use the IBM Cognos Configuration GUI to configure a connection between the Cognos software and an SMTP server. When you configure the connection to the SMTP server, you can send the custom and predefined reports that you create as email attachments.

1. Open the IBM Cognos Configuration GUI.
2. On the Explorer pane, click **Data Access**.
3. Click **Notification**, and then enter values in the following fields:

### SMTP Mail Server

Enter the host name and port number of the mail server.

### Account and password

Leave the field blank if authentication is not required to log on to the mail server. If authentication is required to log on to the mail server, click the **Edit** icon. Enter a valid user ID and password, and then click **OK**.

### Default sender

Enter the email address of the sender.

4. To ensure that the connection to the mail server is correctly configured, right-click **Notification**, and then click **Test**.
5. Click **Save**.
6. Restart the Jazz for Service Management server. To do this, complete the following steps:
  - a. Open a command-line interface.
  - b. Stop the Jazz for Service Management server:

### Windows

```
JazzSM_install_directory\profile\bin\stopServer.bat server_name  
-username JazzSM_username -password JazzSM_user_password
```

### AIX or Linux

```
JazzSM_install_directory/profile/bin/stopServer.sh server_name  
-username JazzSM_username -password JazzSM_user_password
```

- c. Start the Jazz for Service Management server:

### Windows

```
JazzSM_install_directory\profile\bin\startServer.bat server_name
```

### AIX or Linux

```
JazzSM_install_directory/profile/bin/startServer.sh server_name
```

**Related tasks:**

“Starting the Cognos Configuration GUI” on page 500  
You can use the IBM Cognos Configuration GUI to configure Cognos reporting software components. For example, you can configure an archive location for the output of reports, and configure a connection between the Cognos software and an SMTP server.

## Running reports by using the `trcmd` command

Use the `trcmd` command to run IBM Cognos reports and save the report output to a specific location in a file system.

Tivoli Common Reporting includes the Cognos reporting software. You must run the `trcmd` command from the `JazzSM_install_directory/reporting/bin` directory on the system where Tivoli Common Reporting is installed.

Each time that you use the `trcmd` command, you must supply a valid user ID and password. When you issue a command, use the `-user` and `-password` options, as in the following example:

```
trcmd -list -reports -user JazzSM_username -password JazzSM_user_password
```

Before you can run a report by using the `trcmd` command, you must first obtain the following information:

- The path to the report.
- The parameters that the report requires.
- The names of the systems that you want in your report.

Then, you run the report by using the `trcmd` command. Optionally, you can run the report and specify the location for the report output.

1. To get the path to the report, use the `-list` option with the `trcmd` command. For example, on the command line, issue the following command:

```
trcmd -list -reports -user JazzSM_username -password JazzSM_user_password
```

Information about the reports, including the path to the report, is displayed in the command window. For example, information similar to the following sample is displayed:

```
/content/folder[@name='IBM Tivoli Storage Productivity Center Predefined Reports']/folder[@name='Storage Systems']/report[@name='Performance of One Storage System']  
/content/folder[@name='IBM Tivoli Storage Productivity Center Predefined Reports']/folder[@name='Storage Systems']/report[@name='Storage Systems Capacity']  
/content/folder[@name='IBM Tivoli Storage Productivity Center Predefined Reports']/folder[@name='Storage Systems']/report[@name='Storage Systems Historical Capacity']
```

In this example, the path for the Storage Systems Capacity report is `/content/folder[@name='IBM Tivoli Storage Productivity Center Predefined Reports']/folder[@name='Storage Systems']/report[@name='Storage Systems Capacity']`.

2. To see what parameters the report requires, you must run the report without parameters. The output from the command lists the parameters that are required. To run the report, use the `-run` option with the `trcmd` command. For example, on the command line, issue the following command:

```
trcmd -run -report "/content/folder[@name='IBM Tivoli Storage Productivity Center Predefined Reports']/folder[@name='Storage Systems']/report[@name='Storage Systems Capacity']"  
-user JazzSM_username -password JazzSM_user_password
```

The required parameters are displayed in the command window. For example, information similar to the following sample is displayed:

```
CTGTRQ091E The report has required parameters [Storage System,
Storage System Internal Configuration]. Use -parameters flag to specify
the parameters.
```

In this example, the parameters Storage System and Storage System Internal Configuration are required.

3. To get the names of the systems that you want in your report, click the report in the Cognos reporting GUI. A page is displayed which lists the storage systems on which you can run the report.
4. Run the report with the parameters that you require. Use the **-run** and **-parameters** options with the **trcmd** command. For example, on the command line, issue the following command:

```
trcmd -run -report "/content/folder[@name='IBM Tivoli Storage
Productivity Center Predefined Reports']/folder[@name='Storage
Systems']/report[@name='Storage Systems Capacity']" -parameters "Storage
System=SVC-1234-Example" "Storage System=Storwise V7000-5678-1234-Example"
"Storage System Internal Configuration=All Configurations"
-user JazzSM_username -password JazzSM_user_password
```

In this example, the Storage Systems Capacity report runs on the SVC-1234-Example and Storwise V7000-5678-1234-Example systems, regardless of whether the systems are storage systems or storage virtualizers.

5. Optional: To run a report and specify the location for the report output, use the **-distribute**, **-format**, and **-location** options. For example, on the command line, issue the following command:

```
trcmd -distribute -report "/content/folder[@name='IBM Tivoli Storage
Productivity Center Predefined Reports']/folder[@name='Storage
Systems']/report[@name='Storage Systems Capacity']" -format PDF -location
E:\tpcreports\sales\Storage_Systems_Capacity.pdf -parameters "Storage
System=SVC-1234-Example" "Storage System=Storwise V7000-5678-1234-Example"
"Storage System Internal Configuration=All Configurations"
-user JazzSM_username -password JazzSM_user_password
```

In this example, the report is saved in PDF format with the file name Storage\_Systems\_Capacity.pdf in the E:\tpcreports\sales\ directory.

## Specifying locations for saving reports

You can configure reports to save the report output to your local file system. You specify a root directory in the file system in which the output of reports can be saved in the IBM Cognos Configuration GUI. You specify the locations in which you want to save the output of reports. Then, you can configure individual reports to save the report output to a location in the file system. For example, you can schedule a report and configure the report output to be saved to a location like E:\tpcreports\sales or /usr/tpcreports/sales.

### Specifying a root directory for saving reports

You can specify a root directory in the file system in which the output of reports can be saved. To do this, you configure an archive location in IBM Cognos Configuration.

1. Open the IBM Cognos Configuration GUI.
2. Select **Actions > Edit Global Configuration**.
3. Click the **General** tab.

4. In the **Archive Location File System Root** field, enter a file URI. For example, enter `file:///e:/tpcreports` or `file:///usr/tpcreports`.
5. Click **OK**.
6. Select **Actions > Restart**.

**Related tasks:**

“Starting the Cognos Configuration GUI” on page 500

You can use the IBM Cognos Configuration GUI to configure Cognos reporting software components. For example, you can configure an archive location for the output of reports, and configure a connection between the Cognos software and an SMTP server.

## Specifying file system locations for reports

You can specify the locations in which you want to save the output of reports in the IBM Cognos reporting software.

You must configure the root directory of the locations that you specify in the Cognos software. The Cognos software is included in Tivoli Common Reporting, so you must stop and restart Tivoli Common Reporting after you specify the locations.

1. In Tivoli Storage Productivity Center, select **Cognos > View predefined reports**.
2. From the **Launch** menu in the Cognos software, click **Administration**.
3. Click the **Configuration** tab.
4. Click **Dispatchers and Services**.

5. Click the **Define File System Locations** icon .

6. On the Define file system locations page, click **New**.

7. Specify a name and the directory for the location on the file system, and then click **Finish**. For example, enter `sales`. The directory is created the first time that the output of a report for the directory is created. Repeat this step to configure more file system locations.

8. To stop Tivoli Common Reporting, you must stop the Jazz for Service Management server. Follow the instructions for your operating system:

- Windows

```
c:\program files\IBM\JazzSM\profiles\bin\stopServer.bat server_name
-username JazzSM_username -password JazzSM_password
```

- AIX or Linux

```
/opt/IBM/JazzSM/profile/bin/stopServer.sh server_name
-username JazzSM_username -password JazzSM_password
```

9. To restart Tivoli Common Reporting, you must restart the Jazz for Service Management server. Follow the instructions for your operating system:

- Windows

```
c:\program files\IBM\JazzSM\profiles\bin\startServer.bat server_name
```


- AIX or Linux

```
/opt/IBM/JazzSM/profile/bin/startServer.sh server_name
```

## Specifying a report to save the report output to a file system

You can configure a report to save the report output to a location in a file system. You can also specify the file name for the report output, and whether to append a date and time to the file name.

You must configure the locations in which you want to save the output of reports in the IBM Cognos reporting software.


1. In the Cognos software, click the **Schedule** icon  for a report.
2. Select **Override the default values**.
3. In the **Delivery** area, select **Save to the file system**.
4. Optional: Click **Edit the options** to specify the file name to use, and the location of the report on the file system. You can also specify what to do if a file name exists. For example, you can specify to use the report name as the file name. You can click **Make the file names unique and append a timestamp** to add the date and time to the file name.
5. Click **OK**.

## Setting up multiple connections to the TPCDB data source


Tivoli Common Reporting includes the IBM Cognos reporting software. In Tivoli Common Reporting, you can add multiple connections to the Tivoli Storage Productivity Center data source so that you can run reports against multiple Tivoli Storage Productivity Center databases (TPCDBs).

If you create multiple connections, you can run reports for multiple Tivoli Storage Productivity Center servers from one Tivoli Common Reporting server. You can also develop and test reports on multiple TPCDB databases in your storage environment.

After you add multiple connections to the TPCDB data source, when you run a report you select the TPCDB database to connect to.

1. On the Tivoli Common Reporting server, add the remote TPCDB database to the system database catalog, and assign an alias to the database. For example, you can catalog the remote TPCDB database as TPCSRV1.
2. In the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.
3. Click **Launch > IBM Cognos Administration**, and click the **Configuration** tab.
4. Click the **TPCDB** data source, and click the **New Connection** icon .
5. Enter the alias that you created in step 1 in the **Name** field, and then click **Next**.
6. From the **Type** list, select **IBM DB2**, and then click **Next**.
7. Enter the alias that you created in step 1 in the **DB2 database name** field.
8. Click **Signons**, select **Password**, and then enter a user ID and password for the remote database.
9. Click **Test the connection**, and then click **Test**. Verify that the status of the connection is Succeeded.
10. Click **Next**, and then click **Next** again.
11. Click **Finish**.
12. If you scheduled reports, modify the parameters of all scheduled reports to specify the database against which to run the report.

### Related information:

 [Installing Tivoli Common Reporting on a remote server](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_t_install_tcr_remote_server.html)  
 For more information about how to install and configure Tivoli Common Reporting on a remote server, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_t\\_install\\_tcr\\_remote\\_server.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_t_install_tcr_remote_server.html).

## Schedule Management



For more information about report schedules, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cra.10.2.0.doc/c\\_view\\_manage\\_schedules.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cra.10.2.0.doc/c_view_manage_schedules.html).

## Customizing the logo and title of reports

You can change the logo that is displayed in all Tivoli Storage Productivity Center reports and the title of individual reports. Before you modify predefined reports, create backup copies of them.



### Creating a folder in the Cognos software

You can create a folder to store backup copies of the predefined reports that you want to modify or for custom reports that you create.

1. In the IBM Cognos reporting software, click the **Home** icon .
2. Click the tab where you want to create the folder. For example, click **Public Folders**.
3. Click the **New Folder** icon .
4. Specify a name for the folder, and then click **Finish**. For example, enter Backups as the name of the folder.

### Creating a backup copy of a report

Before you modify a predefined report, create a backup copy of the report.

1. In the IBM Cognos reporting software, create a folder in which to store the backup copy.
2. Navigate to the report of which you want to create a backup copy.
3. Click the check box that is next to the title of the report.
4. Click the **Copy** icon .
5. Navigate to the folder that you created in step 1.
6. Click the **Paste** icon .

#### Related tasks:

“Creating a folder in the Cognos software”

You can create a folder to store backup copies of the predefined reports that you want to modify or for custom reports that you create.

### Changing the logo for all reports

You can change the logo that is displayed in the output of all reports. The default logo is the IBM Tivoli logo. You can change the image to your company logo or to another image.


Make sure that your logo is approximately the same size as the IBM Tivoli logo that is in all predefined reports. The image file for the IBM Tivoli logo is at the following location:

```
JazzSM_installation_directory/IBM/JazzSM/profile/installedApps/  
JazzSMNode01Cell/IBM Cognos.ear/p2pd.war/tivoli/tcr_common/images/tivoli.gif
```

1. To store images for your company, create a directory, for example, /myCompany/images, at the following location:  

```
JazzSM_installation_directory/IBM/JazzSM/profile/installedApps/  
JazzSMNode01Cell/IBM Cognos.ear/p2pd.war
```
2. Copy the logo image to the new directory, for example, /myCompany/images/myLogo.jpg.



3. In the IBM Cognos reporting software, click the **Public Folders** tab, and then click **Tivoli Storage Productivity Center Report Layouts**.
4. Create a backup copy of the `report_page_template` file.
5. Click the **Edit with Report Studio** icon  for the `report_page_template` file.
6. Click the IBM Tivoli image in the work area.
7. Double-click the **URL** property in the **Properties** pane.
8. Modify the image URL to refer to your logo image. For example, you can specify the following image URL:  
`../myCompany/images/myLogo.jpg`
9. Select **File > Save**.

**Related tasks:**




“Creating a backup copy of a report” on page 506

Before you modify a predefined report, create a backup copy of the report.

### Modifying the title of a report

You can change the title that is displayed in the output of a report.

You modify the title in Report Studio. You can also change the title of the report by setting a property for the report in the IBM Cognos reporting software. However, if you change this property, then other reports cannot link to the report. That is, if you click a link to the report with the changed property in other reports, the report is not displayed.

1. In the Cognos software, navigate to the location of the report whose title you want to change.
2. Create a backup copy of the report.
3. Click the **Edit with Report Studio** icon  for the report. A message is displayed about the package.
4. Click the **Close** icon  on the message.
5. Click the report header in the work area.
6. Double-click the **Overrides** property in the **Properties** pane.
7. Select the **report\_name** check box, and click **OK**.
8. Click the **Toolbox** icon  in the content pane.
9. Drag a **Text Item** object from the Toolbox and drop the object on the **Drop item to override component child** text in the work area.
10. Type the new title for the report in the Text dialog box, and click **OK**.
11. Select **File > Save**.

To change the title back to the default title, double-click the **Overrides** property, and clear the **report\_name** check box.

**Related tasks:**

“Creating a backup copy of a report” on page 506


Before you modify a predefined report, create a backup copy of the report.


## Creating a search index of reports, report views, and folders

You can create a search index of the reports, report views, and folders in the IBM Cognos reporting software.

Use the search index to search the reports to identify whether the reports contain a property or performance metric.

**Tip:** If there is no search index, you can search only the titles and descriptions of the reports, not the contents.

1. In Tivoli Storage Productivity Center, select **Cognos > View predefined reports**.
2. From the **Launch** menu in Tivoli Common Reporting, click **Administration**.
3. Click the **Index Search** tab.
4. Select the following check boxes in the **Indexable Types** area:
  - **Folder**
  - **Report**
  - **Report view**
5. Click **Save**.
6. Click the **Configuration** tab, and then click **Content Administration**.
7. Click the **New Index Update** icon .
8. Type a name for the index update task, then click **Next**.
9. Select the check box beside each of the items that you want to include in the index update task. For example, select the **Public Folders** check box.
10. Click **Next**.
11. To specify the schedule and scope of the index update task, select one of the following options:
  - **Save and run once:** Specify the scope after you click **Finish**.
  - **Save and schedule:** Specify the schedule and the scope after you click **Finish**.
  - **Save only:** Specify the scope when you run the task.

When you specify the scope, you specify whether to index all items, or only items that changed since the last index was updated.
12. Click **Finish**.
13. To run the index update task immediately, click the **Run with options** icon .

**Related tasks:**

“Searching the predefined reports for a property or performance metric” on page 494

You can search the contents of the predefined reports. You can check which reports contain a property or performance metric. For example, you can list the reports that contain the Storage Pool Name property or the Port Receive I/O Rate performance metric.

**Related information:**

 Create an Index Update Task

For more information about creating a search index, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.svg.ba.cognos.ug\\_cra.10.2.0.doc/t\\_importsearchdata.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.svg.ba.cognos.ug_cra.10.2.0.doc/t_importsearchdata.html).

## Configuring security for the Cognos software with the default authentication configuration

Use WebSphere Application Server to specify which user IDs can log on to the IBM Cognos reporting software. By default, all user IDs have administrative privileges for the Cognos software. In the Cognos software, you can specify which user IDs and groups of user IDs can use reports, report tools, and report administration functions.

You can configure security in the following ways:

- Create user IDs in WebSphere Application Server.
- Create groups of user IDs in the Cognos software.
- Assign administrative privileges to administrator IDs and remove administrative privileges from other user IDs.
- Specify which user IDs or groups can access individual reports or report folders.
- Specify which user IDs or groups can access individual Cognos functions.

The default authentication configuration for Tivoli Storage Productivity Center uses the built-in file repository of the WebSphere Application Server.

### Creating user IDs in WebSphere Application Server

To manage the security of users in the IBM Cognos reporting software, you can create user IDs in WebSphere Application Server.

This task is only required if you use the built-in file repository of the WebSphere Application Server to configure security. This task is not required if you use LDAP or OS authentication methods.

The Cognos software is included in Tivoli Common Reporting. To complete this task, you must open the WebSphere Integrated Solutions Console for the instance of WebSphere Application Server that is running Tivoli Common Reporting.

1. To open the WebSphere Integrated Solutions Console, open a web browser, and open one of the following web addresses:

- `http://hostname:port/ibm/console/logon.jsp`
- `https://hostname:port/ibm/console/logon.jsp`

The host name is the server name or IP address for the server that is running WebSphere Application Server. The port is the port number for the same server.

The port number differs depending on which protocol you used (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center. The default port for http is 16315 and the default port for https is 16316. For example, you might open the following web address:  
`http://storage.example.com:16315/ibm/console/logon.jsp`

2. Log on to the WebSphere Integrated Solutions Console with administrator privileges.

3. To create a new user ID, complete the following steps:

- a. Select **Users and Groups > Manage Users**.
- b. Click **Create**.
- c. Enter the details of the user ID, and then click **Create**.

To create more user IDs, repeat this step.

4. To give administrative privileges to users on WebSphere Application Server, complete the following steps:

- a. Select **Users and Groups > Administrative User Roles**.
- b. Click **Add**.
- c. Search for the user IDs that you want, select the user IDs, and then add the user IDs to the **Mapped to role** list.
- d. Select **Admin Security Manager** and **Administrator** in the roles list.
- e. Click **OK**, and then click **Save**.

## Example: Creating user IDs

Create administrator IDs, normal user IDs, and restricted user IDs.

1. Log on to the WebSphere Integrated Solutions Console.
2. Create two administrator IDs, tcradmin1 and tcradmin2, two normal user IDs, tcnormaluser1, tcnormaluser2, and one restricted user, tcrestricteduser1.
3. Add the roles **Admin Security Manager** and **Administrator** to the two administrator IDs tcradmin1 and tcradmin2.


Create groups of user IDs in the Cognos software.

### Related tasks:

“Creating groups of user IDs in the Cognos software”

To make it easier to manage security, you can create groups of user IDs in the IBM Cognos reporting software.

### Related information:

 Administrative roles and naming service authorization

For more information about administrative roles, see the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEQTP\\_8.5.5/com.ibm.websphere.base.doc/ae/csec\\_adminconsole.html](http://www.ibm.com/support/knowledgecenter/SSEQTP_8.5.5/com.ibm.websphere.base.doc/ae/csec_adminconsole.html).

## Opening the Cognos software

You can use the IBM Cognos reporting software to configure security for Tivoli Storage Productivity Center reports.

Open a web browser, and open the following web address:

`https://servername:port/tarf/servlet/dispatch`

The host name is the server name or IP address for the server that is running the Cognos software. The port is the port number for the same server.


For example, you might open the following web address:

`https://storage.example.com:16311/tarf/servlet/dispatch`

**Tip:** Alternatively, in the Tivoli Storage Productivity Center web-based GUI, select **Cognos > View predefined reports**.

## Creating groups of user IDs in the Cognos software

To make it easier to manage security, you can create groups of user IDs in the IBM Cognos reporting software.

1. Open the Cognos software.
2. Log on to the Cognos software with administrative privileges.
3. From the **Launch** menu, click **Administration**.
4. Click the **Security** tab.
5. Click the **Cognos** directory.
6. To create a group of user IDs, complete the following steps:
  - a. Click the **New Group** icon .
  - b. Enter the details of the group, and then click **Next**.
  - c. Click **Add** on the “Select the members” page.
  - d. Click the **VMMProvider** directory, and after the page refreshes, click **Search**.
  - e. Search for the user ID that you want, and select the check box for the user ID.
  - f. Add the user ID to the **Selected entries** list, and then click **OK**.

g. Click **Finish**.

To create more groups of user IDs, repeat step 6 on page 510.

### **Example: Creating groups of user IDs**

Organize the different types of IDs into groups. It is easier to manage security when you create groups of user IDs. This example uses the user IDs that you created in WebSphere Application Server.

1. Log on to the Cognos software with the tcradmin1 ID.
2. Create a group for administrators, tcradminGroup, and add the user IDs tcradmin1 and tcradmin2 to the group.

**Tip:** If you add the user ID of the administrator that was used to install Jazz for Service Management, the default administrator for Jazz for Service Management becomes an administrator of the Cognos software also.

3. Create a group for normal users, tcnormaluserGroup, and add the user IDs tcnormaluser1 and tcnormaluser2 to the group.
4. Create a group for administrators, tcrrestricteduserGroup, and add the user ID tcrrestricteduser1 to the group.

Restrict administrative privileges in the Cognos software.

#### **Related tasks:**

“Opening the Cognos software” on page 510

You can use the IBM Cognos reporting software to configure security for Tivoli Storage Productivity Center reports.

“Restricting administrative privileges for users in the Cognos software”

By default, all user IDs have administrative privileges. To restrict administrative privileges, you must first add the group of user IDs that you want to have administrative privileges to an administrative role in the IBM Cognos reporting software. Then, you can remove other user IDs from the administrative role.

“Creating user IDs in WebSphere Application Server” on page 509

To manage the security of users in the IBM Cognos reporting software, you can create user IDs in WebSphere Application Server.


### **Restricting administrative privileges for users in the Cognos software**

By default, all user IDs have administrative privileges. To restrict administrative privileges, you must first add the group of user IDs that you want to have administrative privileges to an administrative role in the IBM Cognos reporting software. Then, you can remove other user IDs from the administrative role.

The Cognos software includes predefined security roles. User IDs and groups must be members of security roles to use individual functions. For example, if a user ID is a member of the **System Administrators** role, then that user ID has administrative privileges.

The Cognos software also includes predefined groups. The **Everyone** group includes all user IDs. By default, the **System Administrators** role includes the **Everyone** group. Therefore, by default all user IDs have administrative privileges.

To restrict who has administrative privileges, you must first add the group of user IDs that you want to have administrative privileges to the **System Administrators** role. Then, you must remove the **Everyone** group from the **System Administrators** role.

1. Open the Cognos software.
2. Log on to the Cognos software with administrative privileges.
3. From the **Launch** menu, click **Administration**.
4. Click the **Security** tab.
5. Click the **Properties** icon  for the **System Administrators** role. If the **System Administrators** role is not visible, scroll down.
6. Click the **Members** tab.
7. Click **Add** on the “Set properties” page.
8. Click the **Cognos** directory.
9. Select the check box for the group of user IDs that you want to have administrative privileges, and then add the group to the **Selected entries** list.
10. Click **OK**.
11. Select the check box for the **Everyone** group.
12. Click **Remove** and then click **OK**.

### **Example: Restricting administrative privileges to the users in a group of user IDs**

Assign administrative privileges to one group of user IDs, and remove administrative privileges from other user IDs to control which users can complete administrative tasks.

This example uses the following user IDs and group IDs:

#### **tcradmin1**

This user ID is a member of tcradminGroup. This group has administrative privileges for WebSphere Application Server and the Cognos software.

#### **tcnormaluser1**

This user ID is a member of tcnormalGroup. This group does not have administrative privileges for WebSphere Application Server or the Cognos software.

1. Log on to the Cognos software with the tcradmin1 ID.
2. Add the tcradminGroup group to the **System Administrators** role.
3. Remove the **Everyone** group from the **System Administrators** role.
4. Click **Logoff**.
5. To verify that not all user IDs have administrative privileges, complete the following steps:
  - a. Log on to the Cognos software with the tcnormaluser1 ID.
  - b. Click the **Launch** menu in the Cognos software.

If the **Administration** item is not displayed on the **Launch** menu, not all users have administrative privileges.


#### **Related information:**

##### **Built-in Entries**

For more information about the Everyone group, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cra.10.2.0.doc/c\\_asg\\_builtin\\_entries.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cra.10.2.0.doc/c_asg_builtin_entries.html).

## Specifying access to reports in the Cognos software

You can grant or deny access to reports or to a report folder in the IBM Cognos reporting software. For example, you can prevent a user or group of users from viewing hypervisor reports.

1. Open the Cognos software.
2. Log on to the Cognos software with administrative privileges.
3. Click **IBM Tivoli Productivity Center Predefined Reports**.
4. Click the **Properties** icon  for the report or reports folder to which you want to restrict access.
5. Click the **Permissions** tab.
6. Select the **Override the access permissions acquired from the parent entry** check box.
7. After the page refreshes, click **Add**.
8. Click the **Cognos** directory and after the page refreshes, click **Search**.
9. Search for the user ID or group that you want, and select the check box for the user ID or group.
10. Add the user ID or group to the **Selected entries** list, and then click **OK**.
11. On the **Permissions** tab, select the check box for the user ID or group.
12. Select the **Deny** check box or the **Grant** check box for each of the permissions that you require.
13. Click **OK**.

### Example: Preventing a group of users from viewing hypervisor reports

This example shows how to prevent the users in the `tcnormaluserGroup` group from accessing the reports in the **Hypervisors** folder.


This example uses the following user IDs and group IDs:

#### **tcadmin1**

This user ID is a member of `tcadminGroup`. This group has administrative privileges for WebSphere Application Server and the Cognos software.

#### **tcnormaluser1**

This user ID is a member of `tcnormalGroup`. This group does not have administrative privileges for WebSphere Application Server or the Cognos software.

1. Log on to the Cognos software with the `tcadmin1` ID.
2. Click **IBM Tivoli Productivity Center Predefined Reports**.
3. Click the **Properties** icon  for the **Hypervisors** folder.
4. Click the **Permissions** tab.
5. Select the **Override the access permissions acquired from the parent entry** check box.
6. After the page refreshes, click **Add**.
7. Search the **Cognos** directory for the `tcnormaluserGroup` group, and then add `tcnormaluserGroup` to the **Selected entries** list.
8. On the **Permissions** tab, select the check box for `tcnormaluserGroup`.
9. Select the **Deny** check box for all of the permissions for `tcnormaluserGroup`, and then click **OK**.
10. To verify that user IDs in `tcnormaluserGroup` do not have access to the hypervisor reports, complete the following steps:

- a. Log on to the Cognos software with the user ID and password of a member of `trnormaluserGroup`. For example, log on with the `trnormaluser1` ID.
- b. Click **IBM Tivoli Productivity Center Predefined Reports**.  
If the **Hypervisors** folder is not displayed in the reports list, the `trnormaluser1` ID does not have access to the hypervisor reports.

## Controlling access to Report Studio

You can control access to functions in the IBM Cognos reporting software by using predefined security roles. You assign groups of user IDs to security roles. For example, you can control which users can create or edit reports in Report Studio.

Cognos software includes predefined security roles. User IDs and groups must be members of security roles to use individual functions. For example, if a user ID is a member of the **Authors** role, the user can use the Report Studio tool to create or edit reports.


Cognos software also includes predefined groups. The **Everyone** group includes all authenticated users. By default, the **Authors** role includes the **Everyone** group. Therefore, by default all authenticated users can use Report Studio.

To control access to Report Studio and other Cognos software functions, use the predefined security roles with the groups of user IDs that you create.

For example, to control who can use Report Studio, you take the following general steps:

1. Add the groups of user IDs that you want to use the tools to the **Authors** and **Express Authors** roles.
2. Remove all other user IDs from the **Authors** and **Express Authors** roles. That is, remove the **Everyone** group from the roles.
3. Grant permissions to the **HTML Items in Report** capability to the **Everyone** group.

This procedure describes the general steps in more detail.

1. Open the Cognos software.
2. Log on to the Cognos software with administrator privileges.
3. From the **Launch** menu, click **Administration**.
4. Click the **Security** tab.
5. To assign a group of users to the **Authors** role and remove the **Everyone** group from the role, complete the following steps:
  - a. Click the **Cognos** directory, and then click the **Properties** icon  for the **Authors** role.
  - b. Click the **Members** tab, and then click **Add**.
  - c. Click the **Cognos** directory.
  - d. Select the check box for the group, and then add it to the **Selected entries** list.
  - e. Click **OK**.
  - f. Select the check box for the **Everyone** user.
  - g. Click **Remove**, and then click **OK**.

Repeat this step for the **Express Authors** role.



6. To grant permissions to the **HTML Items in Report** function, complete the following steps:
  - a. On the **Security** tab, click **Capabilities**.
  - b. Click **Report Studio** in the **Capability** list.
  - c. Click the **Actions** icon beside the **HTML Items in Report** list, and then select **Set properties**.
  - d. Click the **Permissions** tab.
  - e. Select the **Override the access permissions acquired from the parent entry** check box.
  - f. Click **Add**.
  - g. Click the **Cognos** directory.
  - h. Select the check box for the **Everyone** group, and then add it to the **Selected entries** list.
  - i. Click **OK**.
  - j. Select the check box for the **Everyone** group, and then elect the **Grant** check box for the **Execute** and **Traverse** permissions.
  - k. Click **OK**.

### Example: Allowing and restricting access to Report Studio

This example shows how to allow access to Report Studio to groups of administrators and other users. The example also shows how to restrict access by some users to these tools.

This example uses the following user IDs and group IDs:

#### **tcradmin1**

This user ID is a member of tcradminGroup. This group has administrative privileges for WebSphere Application Server and the Cognos software.

#### **tcnormaluser1**

This user ID is a member of tcnormalGroup. This group does not have administrative privileges for WebSphere Application Server or the Cognos software.

#### **tcrestricteduser1**

This user ID is a member of tcrestrictedGroup. This group does not have administrative privileges for WebSphere Application Server or the Cognos software.

1. Assign the tcnormalGroup group to the **Authors** role, and remove the **Everyone** group from the role.
2. Assign the tcnormalGroup group to the **Express Authors** role, and remove the **Everyone** group from the role.
3. Assign the tcradminGroup group to the **Authors** and **Express Authors** roles.
4. For the **HTML Items in Report** capability, grant **Execute** and **Traverse** permissions to the **Everyone** group.
5. To verify that users who are in tcrestrictedGroup cannot use Report Studio, complete the following steps:
  - a. Log on to Cognos software with the user ID and password of a member of tcrestrictedGroup. For example, log on with the tcrestricteduser1 ID.
  - b. Click the **Launch** menu in the Cognos software.  
If the Report Studio item is not on the **Launch** menu, users who are in tcrestrictedGroup cannot use the tool.

**Related tasks:**

“Opening the Cognos software” on page 510

You can use the IBM Cognos reporting software to configure security for Tivoli Storage Productivity Center reports.

**Related information:**

 Predefined Entries

For more information about predefined groups and roles like Authors, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cra.10.2.0.doc/c\\_predefinedentries.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cra.10.2.0.doc/c_predefinedentries.html).

---

## Storage virtualizer reports that are available when you upgrade

If you upgrade from Tivoli Storage Productivity Center Version 5.1, Version 5.1.1, or Version 5.1.1.1 to Version 5.2 (or later), you might experience problems with reports.

Review the following table of solutions to common problems with reports:

Problem	Solution
In Tivoli Storage Productivity Center Version 5.2 (or later), no predefined reports are designed specifically for storage virtualizers.	Run predefined reports about storage systems, and select <b>Storage Virtualizer</b> in the <b>Storage System Configuration</b> list.
Tivoli Storage Productivity Center Version 5.2 (or later) contains a reports package, Performance. This package replaces the IBM Tivoli Storage Productivity Center Performance package. When you upgrade from Tivoli Storage Productivity Center Version 5.1, Version 5.1.1, or Version 5.1.1.1 to Version 5.2 (or later), you can continue to use the IBM Tivoli Storage Productivity Center Performance package in Version 5.2 (or later). Reports that are designed specifically for storage virtualizers are only available in this package. However, the following reports from Version 5.1, Version 5.1.1, and Version 5.1.1.1 do not function correctly in Version 5.2 (or later): <ul style="list-style-type: none"> <li>• Storage Systems - Most Active Ports Report</li> <li>• Storage System - Port Details Report</li> <li>• Storage Systems - Most Active Arrays Report</li> <li>• Storage System - Array Details Report</li> </ul>	View one of the following predefined reports about storage systems in Version 5.2 (or later): <ul style="list-style-type: none"> <li>• Most Active Ports</li> <li>• Performance of One Port</li> <li>• Most Active RAID Arrays</li> <li>• Performance of One RAID Array</li> </ul>
If you used the IBM Tivoli Storage Productivity Center Performance package in Version 5.1, Version 5.1.1, or Version 5.1.1.1 to create custom reports, you might need to re-create those reports.	If your custom reports use the following data, you must re-create the reports in the Performance package: <ul style="list-style-type: none"> <li>• Performance metrics for ports on DS8000 and DS6000 systems</li> <li>• Performance metrics for RAID arrays</li> </ul>

Problem	Solution
If you used the IBM Tivoli Storage Productivity Center Resources package in Version 5.1, Version 5.1.1, or Version 5.1.1.1 to create custom reports, you might need to re-create those reports.	If your custom reports use data about VMware file systems, you must re-create the reports in the Capacity and Relationships package.
In Tivoli Storage Productivity Center Version 5.2 (or later), the predefined reports about storage tiering are not in the same folder as Version 5.1, Version 5.1.1, or Version 5.1.1.1.	In Tivoli Storage Productivity Center Version 5.2 (or later), go to the <b>Public Folders &gt; IBM Tivoli Productivity Center Packages &gt; Storage Tiering</b> folder.

## Storage virtualizer resource reports

Storage virtualizer resource reports show the space statistics for storage virtualizers. The reports also show the space statistics for local disks, managed disks, pools, and volumes on storage virtualizers.

### Running the Storage Virtualizers - Space Report

To see information about the capacity of pools, managed disks (MDisks), and volumes on storage virtualizers, run the Storage Virtualizers - Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers in the list are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Resources**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizers - Space Report**.
4. Optional: Select one or more storage virtualizers.
5. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Storage Virtualizers - Space Report”

The Storage Virtualizers - Space Report shows the space statistics for pools, managed disks (MDisks), disks, volumes, and virtual volumes on storage virtualizers.

### Storage Virtualizers - Space Report

The Storage Virtualizers - Space Report shows the space statistics for pools, managed disks (MDisks), disks, volumes, and virtual volumes on storage virtualizers.

### Charts

The bar chart shows the used space and available space for pools on storage virtualizers.

## Report output

For each storage virtualizer, the following information is provided:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Storage Virtualizer Location

The physical location of a storage virtualizer. The location is defined when a storage virtualizer is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage virtualizer in the **Properties** pane for the storage system in the web-based GUI.

### Storage Virtualizer Total Pool Capacity (GiB)

The total amount of storage space in pools on the storage virtualizer.

### Storage Virtualizer Pool Used Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

### Storage Virtualizer Pool Available Space (GiB)

The amount of space in a pool that is not reserved for volumes.

### Storage Virtualizer Physical Allocation Percentage

The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

### Storage Virtualizer Virtual Allocation Percentage

The percentage of physical space in storage virtualizer pools that is

committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

#### **Storage Virtualizer Shortfall Percentage**

The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB  $\div$  100 GiB  $\times$  100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

**Storage Virtualizer Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

**Storage Virtualizer MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage Virtualizer MDisk Available Space (GiB)**

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

**Storage Virtualizer Total Volume Capacity (GiB)**

The amount of virtual volume space on all volumes on the storage virtualizer.

**Storage Virtualizer Volume Assigned Space (GiB)**

The amount of virtual volume space that is assigned to the server or storage virtualizer.

**Storage Virtualizer Disk Capacity (GiB)**

The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

**Storage Virtualizer Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Virtualizer Number of Volumes**

The number of volumes on a resource.

**Storage Virtualizer Number of MDisks**

The number of managed disks on a resource.

**Report links**

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click the name of the report. You can open the following related reports:

- The Storage Virtualizer - Pools Space Report
- The Storage Virtualizer - Volumes Space Report
- The Managed Disks Capacity
- The Storage Virtualizer - Local Disks Space Report
- The Storage Virtualizers - Historical Space Report

**Related tasks:**

“Running the Storage Virtualizers - Space Report” on page 517

To see information about the capacity of pools, managed disks (MDisks), and volumes on storage virtualizers, run the Storage Virtualizers - Space Report.

## Running the Managed Disks Capacity report

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Managed Disks Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Managed Disks Capacity report”

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

## Managed Disks Capacity report

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

**Restriction:** This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

## Charts

The bar chart shows the used space and the available space for managed disk on storage systems.

## Report output

For each managed disk on a storage system, the following information is provided:

### Storage System MDisk Name

The name that was assigned to the managed disk on a storage system when it was added to the system.

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System,

the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage System MDisk Used Space (GiB)**

The amount of used space on the managed disk.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage MDisk RAID Level**

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

**Storage System MDisk Type**

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

**Storage System MDisk Mode**

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

**Array** Extents are provided from local disks.

**Managed**

Extents are provided from a back-end storage volume.

**Unmanaged**

The managed disk is not used in the system.

**Storage System MDisk Strip Size (KB)**

The RAID strip size on a managed disk on a storage system.

**Storage System MDisk Is Balanced**

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is Yes, the LUNs are balanced.

**Storage System MDisk Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

**Storage System MDisk Write Verify**

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is Yes, all write operations are verified by a follow-up read operation.

**Related tasks:**

“Running the Managed Disks Capacity report” on page 521

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

**Running the Storage Virtualizer - Pools Space Report**

To see information about the used space and available space in pools on storage virtualizers, run the Storage Virtualizer - Pools Space Report.



To run the report with the default settings, click **Finish**. By default all of the storage virtualizers in the list are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Resources**.
2. Click Storage Virtualizer.
3. Click **Storage Virtualizer - Pools Space Report**.
4. Optional: Select one or more storage virtualizers.
5. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizer - Pools Space Report”

The Storage Virtualizer - Pools Space Report shows the space statistics for pools on storage virtualizers.

## **Storage Virtualizer - Pools Space Report**

The Storage Virtualizer - Pools Space Report shows the space statistics for pools on storage virtualizers.

### **Charts**

The bar chart shows the used space and available space for pools on storage virtualizers.

### **Report output**

**Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer Pool Capacity (GiB)**

The total amount of storage space in a pool.

**Storage Virtualizer Pool Used Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

**Storage Virtualizer Pool Available Space (GiB)**

The amount of space in a pool that is not reserved for volumes.

**Storage Virtualizer Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

**Storage Virtualizer Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Virtualizer Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Virtualizer Pool Extent Size (MiB)**

The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool Number of Storage Volumes**

The number of volumes in the storage virtualizer pool.

**Storage Virtualizer Pool Number of MDisks**

The number of managed disks on a resource.

**Storage Virtualizer Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Virtualizer Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Virtualizer Pool Physical Allocation Percentage**

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Virtualizer Pool Virtual Allocation Percentage**

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

**Storage Virtualizer Pool Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{total volume capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Virtualizer Pool Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Virtualizer Pool Is Solid State**

Shows whether there are solid-state drives in the pool. This property can contain the following values:

**Mixed** The pool contains both hard disk drives and solid-state drives.

**Non solid state**

The pool contains no solid-state drives.

**Solid state**

The pool contains a least one solid-state drive.

**Storage Virtualizer Pool Warning Level**

The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Report links**

To open related reports, click the name of the storage pool and then click the name of the report. You can open the following related reports:

- The Storage Virtualizer - Volumes Space Report
- The Managed Disks Capacity
- The Storage Virtualizer - Pools Historical Space Report

**Related tasks:**

“Running the Storage Virtualizer - Pools Space Report” on page 522

To see information about the used space and available space in pools on storage virtualizers, run the Storage Virtualizer - Pools Space Report.

**Running the Storage Virtualizer - Volumes Space Report**

To see information about the capacity of volumes on storage virtualizers, run the Storage Virtualizer - Volumes Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers in the list are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Resources**.
2. Click Storage Virtualizer.
3. Click the **Storage Virtualizer - Volumes Space Report** link.
4. Optional: Select one or more storage virtualizers.
5. Optional: Select the type of volume.

Option	Description
All	Select <b>All</b> , the default, to view information about all of the volumes on the selected storage virtualizers.
Primary	Select <b>Primary</b> to view information about primary volumes on the selected storage virtualizers.
Secondary	Select <b>Secondary</b> to view information about secondary volumes on the selected storage virtualizers.

6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizer - Volumes Space Report”

The Storage Virtualizer - Volumes Space Report shows the space statistics for volumes on storage virtualizers.

## Storage Virtualizer - Volumes Space Report

The Storage Virtualizer - Volumes Space Report shows the space statistics for volumes on storage virtualizers.

### Charts

The bar chart shows the used space for volumes on storage virtualizers.

### Report output

For each volume on a storage virtualizer, the following information is provided:

**Storage Virtualizer Volume Name**

The name that was assigned to the storage virtualizer volume when it was added to the system.

**Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer Volume Copy ID**

The identifier of a volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Volume Copy Is Primary**

Shows whether the volume copy is the primary volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Virtualizer Volume Capacity (GiB)**

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

**Storage Virtualizer Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Virtualizer Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

**Storage Virtualizer Volume Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

#### **Storage Virtualizer Volume Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable volume space} \div \text{unallocated volume space}) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$$(\text{100 GiB unallocated volume space} - \text{70 GiB remaining pool space}) \div \text{100 GiB unallocated volume space} \times 100$$

This value is only available for pools with thin-provisioned volumes.

#### **Storage Virtualizer Volume Used Allocated Space Percentage**

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 volumes.

#### **Storage Virtualizer Volume Is Assigned**

Shows whether the volume is assigned to a server. If this value is Yes, the volume is assigned to a server.

#### **Storage Virtualizer Volume Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

#### **Storage Virtualizer Volume Warning Level**

The percentage of used capacity of the volume at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Volume Type**

The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Volume Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Volume RAID Level**

The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

## **Related links**

To open the related report, click the name of the volume, click **Go To**, and then click **Storage Virtualizer - Volumes Historical Space Report**.

### **Related tasks:**

“Running the Storage Virtualizer - Volumes Space Report” on page 526

To see information about the capacity of volumes on storage virtualizers, run the Storage Virtualizer - Volumes Space Report.

## **Storage virtualizer performance reports**

Storage virtualizer performance reports show the most active storage virtualizers, and the most active storage resources on storage virtualizers.

### **Running the Storage Virtualizers - Most Active IO Groups Report**

To see which I/O groups are most active on your storage virtualizers, run the Storage Virtualizers - Most Active IO Groups Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizers - Most Active IO Groups Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.



**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizers - Most Active IO Groups Report”

The Storage Virtualizers - Most Active IO Groups Report shows the performance statistics for the most active I/O groups on storage virtualizers. Use the report to monitor the performance of I/O groups over a specified period.

### **Storage Virtualizers - Most Active IO Groups Report**

The Storage Virtualizers - Most Active IO Groups Report shows the performance statistics for the most active I/O groups on storage virtualizers. Use the report to monitor the performance of I/O groups over a specified period.

### **Charts**

The chart shows the most active I/O groups on storage virtualizers. By default, I/O groups are sorted by the total overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided for I/O groups is as follows:

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer I/O Group Name**

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Total Normal I/O Rate (ops/s)**

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

### **Related reports**

To open related reports, click the name of the resource or storage device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - IO Group Details Report

**Related tasks:**

“Running the Storage Virtualizers - Most Active IO Groups Report” on page 530  
 To see which I/O groups are most active on your storage virtualizers, run the Storage Virtualizers - Most Active IO Groups Report.

**Running the Storage Virtualizers - Most Active MDisks Report**

To see which managed disks (MDisks) are most active on your storage virtualizers, run the Storage Virtualizers - Most Active MDisks Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click Storage Virtualizer.
3. Click **Storage Virtualizers - Most Active MDisks Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizers - Most Active MDisks Report”

The Storage Virtualizers - Most Active MDisks Report shows the performance statistics for the most active managed disks (MDisks) on storage virtualizers. Use the report to monitor the performance of MDisks over a specified period.

**Storage Virtualizers - Most Active MDisks Report**

The Storage Virtualizers - Most Active MDisks Report shows the performance statistics for the most active managed disks (MDisks) on storage virtualizers. Use the report to monitor the performance of MDisks over a specified period.

**Charts**

The chart shows the most active MDisks on storage virtualizers. By default, MDisks are sorted by the total back-end I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

## Report output

The information that is provided for MDisks is as follows:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

### Storage Virtualizer MDisk Name

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

### Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

## Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Pool Details Report
- Storage Virtualizer - MDisk Details Report
- Storage Resource Relationships Summary

### Related tasks:

“Running the Storage Virtualizers - Most Active MDisks Report” on page 532  
To see which managed disks (MDisks) are most active on your storage virtualizers, run the Storage Virtualizers - Most Active MDisks Report.

## Running the Storage Virtualizers - Most Active Nodes Report

To see which nodes are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Nodes Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.

3. Click **Storage Virtualizers - Most Active Nodes Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizers - Most Active Nodes Report”

The Storage Virtualizers - Most Active Nodes Report shows the performance statistics for the most active nodes on storage virtualizers. Use the report to monitor the performance of nodes over a specified period.

## **Storage Virtualizers - Most Active Nodes Report**

The Storage Virtualizers - Most Active Nodes Report shows the performance statistics for the most active nodes on storage virtualizers. Use the report to monitor the performance of nodes over a specified period.

### **Charts**

The chart shows the most active nodes on storage virtualizers. By default, nodes are sorted by the total overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided for nodes is as follows:

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer I/O Group Name**

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

**Storage Virtualizer Node Name**

The name that was assigned to the storage virtualizer node when it was added to the system.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Related reports**

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - IO Group Details Report
- Storage Virtualizer - Node Details Report

**Related tasks:**

“Running the Storage Virtualizers - Most Active Nodes Report” on page 533

To see which nodes are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Nodes Report.

**Running the Storage Virtualizers - Most Active Pools Report**

To see which pools are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Pools Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click Storage Virtualizer.
3. Click **Storage Virtualizers - Most Active Pools Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizers - Most Active Pools Report” on page 536

The Storage Virtualizers - Most Active Pools Report shows the performance statistics for the most active pools on storage virtualizers. Use the report to monitor the performance of pools over a specified period.

## Storage Virtualizers - Most Active Pools Report

The Storage Virtualizers - Most Active Pools Report shows the performance statistics for the most active pools on storage virtualizers. Use the report to monitor the performance of pools over a specified period.

### Charts

The chart shows the most active pools on storage virtualizers. By default, pools are sorted by the total overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

The information that is provided for pools is as follows:

#### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

#### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Pool Details Report
- Storage Virtualizers - Most Active MDisks Report
- Storage Virtualizers - Most Active Volumes Report
- Storage Resource Relationships Summary

#### Related tasks:

“Running the Storage Virtualizers - Most Active Pools Report” on page 535  
To see which pools are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Pools Report.

### Running the Storage Virtualizers - Most Active Volumes Report

To see which volumes are most active on storage virtualizers, run the Storage Virtualizers - Most Active Volumes Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizers - Most Active Volumes Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizers - Most Active Volumes Report”

The Storage Virtualizers - Most Active Volumes Report shows the performance statistics for the most active volumes on storage virtualizers. Use the report to monitor the performance of volumes over a specified period.

### **Storage Virtualizers - Most Active Volumes Report**

The Storage Virtualizers - Most Active Volumes Report shows the performance statistics for the most active volumes on storage virtualizers. Use the report to monitor the performance of volumes over a specified period.

### **Charts**

The chart shows the most active volumes on storage virtualizers. By default, volumes are sorted by the total overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided for volumes is as follows:

#### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### **Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

#### **Storage Virtualizer I/O Group Name**

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

**Storage Virtualizer Volume Name**

The name that was assigned to the storage virtualizer volume when it was added to the system.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Related reports**

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Pool Details Report
- Storage Virtualizer - IO Group Details Report
- Storage Virtualizer - Volume Details Report
- Storage Resource Relationships Summary

**Related tasks:**

“Running the Storage Virtualizers - Most Active Volumes Report” on page 536  
To see which volumes are most active on storage virtualizers, run the Storage Virtualizers - Most Active Volumes Report.

**Running the Storage Virtualizers - Most Active Ports Report**

To see which ports are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Ports Report.

To run the report with the default settings, click **Finish**.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizers - Most Active Ports Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.



If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizers - Most Active Ports Report”

The Storage Virtualizers - Most Active Ports Report shows the performance statistics for the most active ports on storage virtualizers. Use the report to monitor the performance of ports over a specified period.

## **Storage Virtualizers - Most Active Ports Report**

The Storage Virtualizers - Most Active Ports Report shows the performance statistics for the most active ports on storage virtualizers. Use the report to monitor the performance of ports over a specified period.

### **Charts**

The chart shows the most active ports on storage virtualizers. By default, ports are sorted by the total port I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided for ports is as follows:

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer Port Name**

An alphanumeric string that identifies a storage virtualizer port.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

### **Related reports**

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Port Details Report

**Related tasks:**

“Running the Storage Virtualizers - Most Active Ports Report” on page 538

To see which ports are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Ports Report.

## Storage virtualizer performance details reports

Storage virtualizer reports about performance details show information about the performance of storage virtualizers, and information about the performance of the resources related to storage virtualizers.

### Running the Storage Virtualizer - Details Report

To see information about the performance of storage virtualizers, run the Storage Virtualizer - Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer Performance Details Report**.
4. Select a storage virtualizer.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Storage Virtualizer - Details Report”

The Storage Virtualizer - Details Report shows the performance statistics for a storage virtualizer. Use the report to monitor the performance of a storage virtualizer over a specified period.

### Storage Virtualizer - Details Report

The Storage Virtualizer - Details Report shows the performance statistics for a storage virtualizer. Use the report to monitor the performance of a storage virtualizer over a specified period.

#### Charts

The charts show the following performance statistics for the storage virtualizer:

- Total overall rates for I/O operations, and the total overall response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

#### Report output

The information that is provided for a storage virtualizer is as follows:

##### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Day or Hour or Time**

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

## Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizers - Most Active IO Groups Report
- Storage Virtualizers - Most Active MDisks Report
- Storage Virtualizers - Most Active Nodes Report
- Storage Virtualizers - Most Active Pools Report
- Storage Virtualizers - Most Active Ports Report
- Storage Virtualizers - Most Active Volumes Report

### Related tasks:

“Running the Storage Virtualizer - Details Report” on page 540

To see information about the performance of storage virtualizers, run the Storage Virtualizer - Details Report.

## Running the Storage Virtualizer - IO Group Details Report

To see information about the performance of an I/O group on a storage virtualizer, run the Storage Virtualizer - IO Group Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - IO Group Details Report**.
4. Select a storage virtualizer, and then select an I/O group.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Storage Virtualizer - IO Group Details Report”

The Storage Virtualizer - IO Group Details Report shows the performance statistics for an I/O group on a storage virtualizer. Use the report to monitor the performance of an I/O group over a specified period.

## Storage Virtualizer - IO Group Details Report

The Storage Virtualizer - IO Group Details Report shows the performance statistics for an I/O group on a storage virtualizer. Use the report to monitor the performance of an I/O group over a specified period.

## Charts

The charts show the following performance statistics for an I/O group on a storage virtualizer:

- The total overall I/O rates, and the total response times

- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

## Report output

The information that is provided for an I/O group is as follows:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Storage Virtualizer I/O Group Name

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

### Day or Hour or Time

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

### Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations.

### Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

#### **Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

#### **Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

### **Related reports**

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active IO Groups Report**.

#### **Related tasks:**

“Running the Storage Virtualizer - IO Group Details Report” on page 542

To see information about the performance of an I/O group on a storage virtualizer, run the Storage Virtualizer - IO Group Details Report.

### **Running the Storage Virtualizer - MDisk Details Report**

To see information about the performance of a managed disk (MDisk) on a storage virtualizer, run the Storage Virtualizer - MDisk Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - MDisk Details Report**.
4. Select a storage virtualizer, and then select an MDisk.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### **Related reference:**

“Storage Virtualizer - MDisk Details Report”

The Storage Virtualizer - MDisk Details Report shows the performance statistics for a managed disk (MDisk) on a storage virtualizer. Use the report to monitor the performance of an MDisk over a specified period.

### **Storage Virtualizer - MDisk Details Report**

The Storage Virtualizer - MDisk Details Report shows the performance statistics for a managed disk (MDisk) on a storage virtualizer. Use the report to monitor the performance of an MDisk over a specified period.

## Charts

The charts show the following performance statistics for an MDisk on a storage virtualizer:

- The total I/O rates for back-end operations, and the total response times for back-end operations
- The I/O rates for read and write back-end operations, and the response times for read and write back-end operations
- The total data rates for back-end operations, and the data rates for read and write back-end operations
- The overall queue times for back-end operations, and the queue times for read and write back-end operations

## Report output

The information that is provided for an MDisk is as follows:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

### Storage Virtualizer MDisk Name

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

### Day or Hour or Time

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

### Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

### Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Related reports**

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active MDisks Report**.

**Related tasks:**

“Running the Storage Virtualizer - MDisk Details Report” on page 544

To see information about the performance of a managed disk (MDisk) on a storage virtualizer, run the Storage Virtualizer - MDisk Details Report.

**Running the Storage Virtualizer - Node Details Report**

To see information about the performance of a node on a storage virtualizer, run the Storage Virtualizer - Node Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - Node Details Report**.
4. Select a storage virtualizer, and then select a node.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.



## 7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### **Related reference:**

“Storage Virtualizer - Node Details Report”

The Storage Virtualizer - Node Details Report shows the performance statistics for a node on a storage virtualizer. Use the report to monitor the performance of a node over a specified period.

## **Storage Virtualizer - Node Details Report**

The Storage Virtualizer - Node Details Report shows the performance statistics for a node on a storage virtualizer. Use the report to monitor the performance of a node over a specified period.

### **Charts**

The charts show the following performance statistics for a node on a storage virtualizer:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

### **Report output**

The information that is provided for a node is as follows:

#### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### **Storage Virtualizer I/O Group Name**

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

#### **Storage Virtualizer I/O Group Name**

To be provided.

#### **Storage Virtualizer Node Name**

The name that was assigned to the storage virtualizer node when it was added to the system.

#### **Day or Hour or Time**

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

#### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Related reports**

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Nodes Report**.

**Related tasks:**

“Running the Storage Virtualizer - Node Details Report” on page 546

To see information about the performance of a node on a storage virtualizer, run the Storage Virtualizer - Node Details Report.

**Running the Storage Virtualizer - Pool Details Report**

To see information about the performance of a pool on a storage virtualizer, run the Storage Virtualizer - Pool Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.

2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - Pool Details Report**.
4. Select a storage virtualizer, and then select a pool.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Virtualizer - Pool Details Report”

The Storage Virtualizer - Pool Details Report shows the performance statistics for a pool on a storage virtualizer. Use the report to monitor the performance of a pool over a specified period.

## **Storage Virtualizer - Pool Details Report**

The Storage Virtualizer - Pool Details Report shows the performance statistics for a pool on a storage virtualizer. Use the report to monitor the performance of a pool over a specified period.

## **Charts**

The charts show the following performance statistics for a pool on a storage virtualizer:

- The total overall rates for I/O operations, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The total I/O rates for back-end operations, and overall response times for back-end operations
- The I/O rates for back-end read operations and back-end write operations, and the response rates for back-end read and back-end write operations

## **Report output**

The information that is provided for a pool is as follows:

### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### **Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

### **Day or Hour or Time**

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Related reports**

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Pools Report**.

**Related tasks:**

“Running the Storage Virtualizer - Pool Details Report” on page 548  
To see information about the performance of a pool on a storage virtualizer, run the Storage Virtualizer - Pool Details Report.

### **Running the Storage Virtualizer - Volume Details Report**

To see information about the performance of a volume on a storage virtualizer, run the Storage Virtualizer - Volume Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - Volume Details Report**.
4. Select a storage virtualizer.
5. Select a volume. To select a volume, complete the following steps:
  - a. Type a keyword. For example, type lun, to find the names of volumes that start with lun. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type a% to find volume names that begin with 'a' or 'A'. Type %a to find volume names that begin with or contain the letter 'a' or 'A'. You can type % to retrieve all of the volumes on a resource.
  - b. Click **Search**. If the search is successful, the names of the volumes are displayed in the **Choice** list.
  - c. From the list, select a volume.
6. Optional: Select an interval.
7. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### **Related reference:**

“Storage Virtualizer - Volume Details Report”

The Storage Virtualizer - Volume Details Report shows the performance statistics of a volume on a storage virtualizer. Use the report to monitor the performance of a volume over a specified period.

### **Storage Virtualizer - Volume Details Report**

The Storage Virtualizer - Volume Details Report shows the performance statistics of a volume on a storage virtualizer. Use the report to monitor the performance of a volume over a specified period.

#### **Charts**

The charts show the following performance statistics for a volume on a storage virtualizer:

- The total overall rates for I/O operations, and the total response times

- The overall rates for I/O read operations, the overall rates for I/O write operations, the response times for read operations, and the response times for write operations
- The data rates for read operations, the data rates for write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for volumes

## Report output

The information that is provided for a volume is as follows:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Day or Hour or Time

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

### Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

### Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

### Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations.

### Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### **Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### **Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

### **Volume Utilization**

The average percentage of time that the volume is busy.

## **Related reports**

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Volumes Report**.

#### **Related tasks:**

“Running the Storage Virtualizer - Volume Details Report” on page 551

To see information about the performance of a volume on a storage virtualizer, run the Storage Virtualizer - Volume Details Report.

## **Running the Storage Virtualizer - Port Details Report**

To see information about the performance of a port on a storage virtualizer, run the Storage Virtualizer - Port Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - Port Details Report**.
4. Select a storage virtualizer, and then select a port.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### **Related reference:**

“Storage Virtualizer - Port Details Report” on page 554

The Storage Virtualizer - Port Details Report shows the performance statistics for a port on a storage virtualizer. Use the report to monitor the performance of a port over a specified period.

## Storage Virtualizer - Port Details Report

The Storage Virtualizer - Port Details Report shows the performance statistics for a port on a storage virtualizer. Use the report to monitor the performance of a port over a specified period.

### Charts

The charts show the following performance statistics for a port on a storage virtualizer:

- The total I/O rates
- The send I/O rates for ports, and the receive I/O rates for ports
- The send data rates for ports, the receive data rates for ports, and the total data rates for ports
- The overall bandwidth percentage for ports

### Report output

The information that is provided for ports is as follows:

#### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### Storage Virtualizer Port Name

An alphanumeric string that identifies a storage virtualizer port.

#### Day or Hour or Time

The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

#### Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port.

#### Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data.

#### Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

#### Port Send Data Rate (MiB/s)

The average rate at which data is sent through the port. The rate is measured in MiB per second.

#### Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second.

#### Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

#### Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.



## Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Ports Report**.

### Related tasks:

“Running the Storage Virtualizer - Port Details Report” on page 553

To see information about the performance of a port on a storage virtualizer, run the Storage Virtualizer - Port Details Report.

## Predefined reports about historical information

You can run predefined reports about historical information to monitor how space is used on storage systems and storage virtualizers over a specified period.

You can run predefined reports that show how space is used on storage systems and related resources such as pools and volumes.

### Running the Storage Virtualizers - Historical Space Report

To see how space on storage virtualizers is used, run the Storage Virtualizers - Historical Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Historical Information**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizers - Historical Space Report**.
4. Optional: Select one or more storage virtualizers.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Storage Virtualizers - Historical Space Report

The Storage Virtualizers - Historical Space Report shows the usage of space on storage virtualizers over a specified period. Use the report to determine whether additional storage resources are required.

## Charts

The chart shows the following space statistics for the storage virtualizers:

- The capacity of the pools on the storage virtualizers
- The used space in the pools on the storage virtualizers
- The available space in the pools on the storage virtualizers
- The capacity of the virtual volumes on the storage virtualizers
- The capacity of the managed disks (MDisks) on the storage virtualizers

## Report output

The following information is included in reports for storage virtualizers:

### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Date** The date that the data was collected for the report.

### **Storage Virtualizer Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

### **Storage Virtualizer Pool Used Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

### **Storage Virtualizer Pool Available Space (GiB)**

The amount of space in a pool that is not reserved for volumes.

### **Storage Virtualizer Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

**Storage Virtualizer MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage Virtualizer MDisk Available Space (GiB)**

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

**Storage Virtualizer Total Volume Capacity (GiB)**

The amount of virtual volume space on all volumes on the storage virtualizer.

**Storage Virtualizer Volume Assigned Space (GiB)**

The amount of virtual volume space that is assigned to the server or storage virtualizer.

**Storage Virtualizer Disk Capacity (GiB)**

The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

**Storage Virtualizer Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Virtualizer Number of Volumes**

The number of volumes in the storage virtualizer pool.

**Storage Virtualizer Number of MDisks**

The number of managed disks on a resource.

## **Running the Storage Virtualizer - Pools Historical Space Report**

To see how space in pools on a storage virtualizer is used, run the Storage Virtualizer - Pools Historical Space Report.

You can see how space in a pool on a storage virtualizer, or in pools in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Historical Information**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - Pools Historical Space Report**.
4. Required: Select the type of historical space report that you want to create.
5. Optional: You can filter the type of data that you want to see in a report. You can select either of the following options:
  - **Of a Pool in a Storage Virtualizer**
    - a. Click **Next**.
    - b. Select a storage virtualizer, and then select a pool.
  - **Of Pools in a Storage Resource Group**
    - a. Click **Next**.
    - b. Select a storage resource group.
6. Optional: Select an interval.

7. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

## **Storage Virtualizer - Pools Historical Space Report**

The Storage Virtualizer - Pools Historical Space Report shows the usage of space in a pool on a storage virtualizer, or of pools in a storage resource group over a specified period. Use the report to determine whether additional storage resources are required.

### **Charts**

The chart shows the following space statistics for the pool on the storage virtualizer:

- The capacity of the pool
- The used space in the pool
- The available space in the pool
- The virtual volume capacity of the pool
- The real volume capacity of the pool

The chart shows the following space statistics for the pools in the storage resource group:

- The capacity of the pools
- The used space in the pools
- The available space in the pools
- The virtual volume capacity of the pools
- The real volume capacity of the pools

### **Report output**

The following information is included in reports for a pool on a storage virtualizer, or for pools in a storage resource group:

#### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### **Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

**Date** The date that the data was collected for the report.

#### **Storage Virtualizer Pool Capacity (GiB)**

The total amount of storage space in a pool.

**Storage Virtualizer Pool Available Space (GiB)**

The amount of space in a pool that is not reserved for volumes.

**Storage Virtualizer Pool Used Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

**Storage Virtualizer Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Virtualizer Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

**Storage Virtualizer Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Virtualizer Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Virtualizer Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Virtualizer Pool Number of Storage Volumes**

The number of volumes in the storage virtualizer pool.

**Storage Virtualizer Pool Number of MDisks**

The number of managed disks on a resource.

**Running the Storage Virtualizer - Volumes Historical Space Report**

To see how space on volumes on a storage virtualizer is used, run the Storage Virtualizer - Volumes Historical Space Report

You can see how space on a volume on a storage virtualizer, or on volumes in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Historical Information**.
2. Click **Storage Virtualizer**.
3. Click **Storage Virtualizer - Volumes Historical Space Report**.
4. Required: Select the type of historical space report that you want to create.
5. Optional: You can filter the type of data that you want to see in a report. You can select either of the following options:
  - **Of a Volume in a Storage Virtualizer**
    - a. Click **Next**.
    - b. Select a storage virtualizer.
    - c. Type a keyword and click **Search**. For example, type `lun`, to find the names of volumes that start with `lun`. To change the default search option, click **Options**.

**Tip:** Use the percent sign (%) as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
    - d. Select a volume.
  - **Of Volumes in a Storage Resource Group**
    - a. Click **Next**.
    - b. Select a storage resource group.
6. Optional: Select an interval.
7. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of filters.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

## **Storage Virtualizer - Volumes Historical Space Report**

The Storage Virtualizer - Volumes Historical Space Report shows the usage of space on a volume on a storage virtualizer, or on volumes in a storage resource group over a specified period. Use the report to determine whether additional storage resources are required.

### **Charts**

The chart shows the following space statistics for the volume on the storage virtualizer:

- The capacity of the volume
- The used space on the volume
- The real space on the volume

The chart shows the following space statistics for the volumes in the storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The real space on the volumes

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

### **Report output**

The following information is included in reports for a volume on a storage virtualizer, or for volumes in a storage resource group:

#### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### **Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

#### **Storage Virtualizer Volume Name**

The name that was assigned to the storage virtualizer volume when it was added to the system.

**Date** The date that the data was collected for the report.

#### **Storage Virtualizer Volume Capacity (GiB)**

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

### **Storage Virtualizer Volume Overhead**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

### **Storage Virtualizer Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

### **Storage Virtualizer Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

## **Running the Storage Systems Historical Capacity report**

To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Storage Systems Historical Capacity**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.



If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Systems Historical Capacity report”

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

## **Storage Systems Historical Capacity report**

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

### **Charts**

The chart shows the following space statistics for the storage systems:

- The capacity of the pools on the storage systems
- The used space in the pools on the storage systems
- The available space in the pools on the storage systems
- The capacity of the volumes on the storage systems
- The available space on the managed disks on the storage systems
- The capacity of the managed disks on the storage systems

### **Report output**

The following information is included in reports for storage systems:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is *Storage System*, the storage system is configured as back-end storage. If this value is *Storage Virtualizer*, the storage system is configured to virtualize back-end storage.

**Date** The date that the data was collected for the report.

**Storage System Pool Capacity (GiB)**

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

**Storage System Used Pool Space (GiB)**

The amount of space that is in use in all pools on a storage system.

**Storage System Available Pool Space (GiB)**

The amount of unused space in pools that are on the storage system.

**Storage System Volume Capacity (GiB)**

The amount of space on all volumes on the storage system.

**Storage System Number of Volumes**

The number of volumes on a resource.

**Storage System Number of Disks**

The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

**Storage System Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage System Unallocated Disk Space (GiB)**

The amount of disk space that can be added to a pool.

**Storage System Volume Capacity Assigned to MDisks (GiB)**

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

**Storage System Disk Capacity (GiB)**

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Total Disk Capacity (GiB)**

The amount of space on physical disks on a storage system, including spare disks.

**Storage System Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage System Volume Capacity for z/OS (GiB)**

The amount of space on all volumes on the storage system that the z/OS operating system can use.

**Storage System Real Available Pool Space (GiB)**

The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Real Configured Pool Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Related tasks:**

“Running the Storage Systems Historical Capacity report” on page 562  
To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

**Running the Pools Historical Capacity report**

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

You can see how space in a pool on a storage system, or in pools in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Pools Historical Capacity**.
5. Select the type of historical space report that you want to create. You can select either of the following options:
  - **Of a Pool in a Storage System**
    - a. Click **Next**.
    - b. Optional: Select the configuration of the storage system.
    - c. Select a storage system, and then select a pool.
  - **Of Pools in a Storage Resource Group**
    - a. Click **Next**.
    - b. Optional: Select the configuration of the storage system.
    - c. Select a storage resource group.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Pools Historical Capacity report”

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

### **Pools Historical Capacity report**

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

### **Charts**

The chart shows the following space statistics for the pool on a storage system:

- The capacity of the pool
- The used space in the pool
- The available space in the pool
- The virtual volume capacity of the pool
- The used volume space in the pool
- The allocated space in the pool

The chart shows the following space statistics for the pool in the storage resource group:

- The capacity of the pools
- The used space in the pools
- The available space in the pools

- The virtual volume capacity of the pools
- The used volume space in the pools
- The allocated space in the pool

## Report output

The following information is included in reports for a pool on a storage system or for pools in a storage resource group:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is *Storage System*, the storage system is configured as back-end storage. If this value is *Storage Virtualizer*, the storage system is configured to virtualize back-end storage.

### Storage Pool Name

The name that was assigned to the pool when it was added to the system.

**Date** The date that the data was collected for the report.

### Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

### Storage Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{pool capacity} - \text{used space}$$

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

### Storage Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Pool Used Space (GiB)**

The amount of storage space that is used in the pool.

**Storage Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Pool Real Configured Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage Pool Number of MDisks**

The number of managed disks in a storage pool.

**Storage Pool Number of Storage Volumes**

The number of volumes in the storage pool.

**Related tasks:**

“Running the Pools Historical Capacity report” on page 564

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

**Running the Volumes Historical Capacity report**

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

You can see how space on a volume on a storage system, or on volumes in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Volumes Historical Capacity**.
5. Select the type of historical capacity report that you want to create. You can select either of the following options:
  - **Of a Volume in a Storage System**
    - a. Click **Next**.
    - b. Optional: Select the configuration of the storage system.
    - c. Select a storage system.
    - d. Type a keyword and click **Search**. For example, type lun, to find the names of volumes that start with lun. To change the default search option, click **Options**.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Volumes Historical Capacity report”

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

**Volumes Historical Capacity report**

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

**Charts**

The chart shows the following space statistics for the volume on a storage system:

- The capacity of the volume
- The used space on the volume
- The space on the volume that is used for system management

- The allocated space on the volume

The chart shows the following space statistics for the volumes in a storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The space on the volumes that is used for system management
- The allocated space on the volume

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

## Report output

The following information is included in reports for a volume on a storage system, or for volumes in a storage resource group:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### Storage Pool Name

The name that was assigned to the pool when it was added to the system.

### Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

**Date** The date that the data was collected for the report.

### Storage Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

### Storage Volume Used Space (GiB)

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Related tasks:**

“Running the Volumes Historical Capacity report” on page 567

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

## Storage virtualizer custom reports about performance

You can use IBM Tivoli Common Reporting to create custom performance reports about storage virtualizers. Performance reports can contain detailed information about the performance of monitored storage virtualizers, and some of the properties of those storage virtualizers.

You can create performance reports for storage virtualizers and their components. You can add performance metrics for volumes, back-end arrays, caches, ports, port errors, and other metrics to performance reports. For example, you can add key metrics to a report to create and generate a performance chart for a specific volume of a storage virtualizer. All the performance metrics from the stand-alone GUI are available in the new reporting interface.

**Data for storage virtualizers in performance reports**

You can include general information, capacity data, properties, and other information about storage virtualizers in performance reports.

**Information about storage virtualizers**

You can create performance reports that include the following information:

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage Virtualizer Total Pool Capacity (GiB)**

The total amount of storage space in pools on the storage virtualizer.



**Storage Virtualizer Pool Used Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

**Storage Virtualizer Pool Available Space (GiB)**

The amount of space in a pool that is not reserved for volumes.

**Storage Virtualizer Total Volume Capacity (GiB)**

The amount of virtual volume space on all volumes on the storage virtualizer.

**Storage Virtualizer Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

**Storage Virtualizer Volume Assigned Space (GiB)**

The amount of virtual volume space that is assigned to the server or storage virtualizer.

**Storage Virtualizer MDisk Available Space (GiB)**

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

**Storage Virtualizer MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage Virtualizer Disk Capacity (GiB)**

The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

**Storage Virtualizer Number of MDisks**

The number of managed disks on a resource.

**Storage Virtualizer Number of Volumes**

The number of volumes on a resource.

**Storage Virtualizer Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Virtualizer Physical Allocation Percentage**

The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Virtualizer Virtual Allocation Percentage**

The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Shortfall Percentage**

The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

## Component properties

You can create performance reports that include the following information:

### Storage Virtualizer Access Information

The URL from which you can access the storage system or storage virtualizer.

### Storage Virtualizer Alias

Shows an alternative name for a storage virtualizer.

### Storage Virtualizer Code Level

The microcode or firmware level of the storage virtualizer.

### Storage Virtualizer Time Zone

The time zone in which a resource is located.

### Storage Virtualizer Custom Tag 1, 2, and 3

User-defined text that is associated with a storage virtualizer. You can add or edit the custom tags for a storage virtualizer on the **Properties** page of the storage system in the Tivoli Storage Productivity Center web-based GUI.

### Storage Virtualizer User Provided Name

The name that was specified for the storage virtualizer in the storage environment.

### Storage Virtualizer Is Compression Active

Shows whether the compression feature is enabled on the storage virtualizer. If this value is Yes, the compression feature is enabled.

## Status information

You can create performance reports that include the following information:

**Storage Virtualizer Status**

The condition of the resource, for example normal, warning, or error.

**Storage Virtualizer Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Virtualizer Last Data Collection**

The date and time when storage statistics were last collected from the resource.

**Storage Virtualizer Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

**Vendor, model, and device information**

You can create performance reports that include the following information:

**Storage Virtualizer IP Address**

The IP address of the resource.

**Storage Virtualizer Serial Number**

The serial number of the resource.

**Storage Virtualizer Type**

The type of storage virtualizer. For example, the storage virtualizer can be an **IBM Storwize V7000** system, an **IBM SAN Volume Controller** system, or another type of storage virtualizer.

**Storage Virtualizer Model**

The model name or model number of the resource.

**Storage Virtualizer Vendor**

The vendor who supplied the resource.

**Performance metrics for storage virtualizers:**

You can include performance data for ports, port errors, volumes, caches, back-end arrays, and other data for storage virtualizers in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

**Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

**Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

**Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

**Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

**Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

**Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

**Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Local Node Response Time (ms/op)**

The average number of milliseconds to complete a send or receive

operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

**Overall Port to Local Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Remote Node Response Time (ms/op)**

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

**Overall Port to Remote Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

**Port error data**

You can create performance reports that include the following information:

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second

for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Volume data**

You can create performance reports that include the following information:

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.



**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

**Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

**Performance metrics for storage systems that are configured to virtualize storage:**

A storage environment can include a Hitachi Universal Storage Platform V system that is configured to virtualize storage. You can include performance data for volumes, and cache data for these systems in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

## Cache data

You can create performance reports that include the following information:

### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

### Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

## Data for managed disks on storage virtualizers in performance reports

You can include general information, capacity data, properties, and other information about managed disks on storage virtualizers in performance reports.

## Information about managed disks on storage virtualizers

You can create performance reports that include the following information:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

### Storage Virtualizer MDisk Name

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

## Capacity and usage data

You can create performance reports that include the following information:

### Storage Virtualizer MDisk Capacity (GiB)

The amount of storage space on the managed disk.

### Storage Virtualizer MDisk Available Space (GiB)

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

### Storage Virtualizer MDisk Strip Size (KB)

The RAID strip size on a managed disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

## Component properties

You can create performance reports that include the following information:

### Storage Virtualizer MDisk RAID Level

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer MDisk Spare Goal

The number of spare drives that are required to maintain redundancy. Use spare drives to protect against drive failures in the array on a managed disk on a storage virtualizer.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer MDisk Type

The type of managed disk on a storage virtualizer. For example, the storage virtualizer can be a local or virtual managed disk.

### Storage Virtualizer MDisk Mode

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

**Array** Extents are provided from local disks.

#### **Managed**

Extents are provided from a back-end storage volume.

#### **Unmanaged**

The managed disk is not used in the system.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer MDisk Is Balanced

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is Yes, the LUNs are balanced.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer MDisk Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer MDisk Write Verify

Shows whether all write operations on a managed disk on a storage virtualizer are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is Yes, all write operations are verified by a follow-up read operation.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

## Status information

You can create performance reports that include the following information:

### Storage Virtualizer MDisk Status

The condition of the resource, for example normal, warning, or error.

### Storage Virtualizer MDisk Native Status

Shows the level of access of nodes in the system to a managed disk or volume on a storage virtualizer. The level of access can be online, offline, degraded, or excluded.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

## Performance metrics for managed disks on storage virtualizers:

You can include performance data for back-end arrays on storage-virtualizer managed disks in performance reports.

### Back-end array data

You can create performance reports that include the following information:

#### Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

#### Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

#### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

#### Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

#### Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

#### Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

#### Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

#### Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.



**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Data for storage virtualizer volumes in performance reports**

You can include general information, capacity data, properties, and other information about storage virtualizer volumes in performance reports.

**Information about storage virtualizer volumes**

You can create performance reports that include the following information:

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Virtualizer I/O Group Name**

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

**Storage Virtualizer Volume Name**

The name that was assigned to the storage virtualizer volume when it was added to the system.

**Storage Virtualizer Volume Assigned Host Alias**

The user-defined name that is assigned to the port on the server. Volumes are associated with the name of the port on the server.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage Virtualizer Volume Capacity (GiB)**

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

**Storage Virtualizer Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

**Storage Virtualizer Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Virtualizer Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Virtualizer Volume Used Allocated Space Percentage**

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 volumes.

#### **Storage Virtualizer Volume Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine the value:

$$\text{capacity} - \text{allocated space}$$

#### **Storage Virtualizer Volume Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

#### **Storage Virtualizer Volume Unused Space (GiB)**

The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 pools.

#### **Storage Virtualizer Volume Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable volume space} \div \text{unallocated volume space}) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$(100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space}) \div 100 \text{ GiB unallocated volume space} \times 100$

This value is only available for pools with thin-provisioned volumes.

**Storage Virtualizer Volume Unallocatable Space (GiB)**

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

**Storage Virtualizer Volume Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

**Component properties**

You can create performance reports that include the following information:

**Storage Virtualizer Volume Type**

The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Volume WWN**

The worldwide name of the volume.

**Storage Virtualizer Volume Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

**Storage Virtualizer Volume Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

**Storage Virtualizer Volume Is Assigned**

Shows whether the volume is assigned to a server. If this value is Yes, the volume is assigned to a server.

**Storage Virtualizer Volume Is Auto Expand Enabled**

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is Yes, the volume or volume copy can allocate new extents from a pool automatically.

**Storage Virtualizer Volume Is Encryptable**

Shows whether the resource can be encrypted. If this value is Yes, the resource can be encrypted.

**Storage Virtualizer Volume Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

**Storage Virtualizer Volume Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Virtualizer Volume Metro Mirror Name**

The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

**Storage Virtualizer Volume Mirror Count**

The number of mirrors that keep a synchronous copy of the resource.

**Storage Virtualizer Volume Preferred Node**

The number of the storage virtualizer node that is used for I/O operations for the storage virtualizer volume.

**Storage Virtualizer Volume Throttle**

The maximum number of commands that the storage virtualizer volume can queue.

**Storage Virtualizer Volume Warning Level**

The percentage of used capacity of the resource at which a warning is generated.

**Storage Virtualizer Volume Copy Relationship**

Shows whether a volume or a copy of a volume is in a FlashCopy relationship or a point-in-time relationship. This property can contain the following values:

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Storage Virtualizer Volume Is Compressed**

Shows whether the storage virtualizer volume is compressed. If this value is Yes, the volume is compressed.

**Storage Virtualizer Volume RAID Level**

The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Status information**

You can create performance reports that include the following information:

**Storage Virtualizer Volume Status**

The condition of the resource, for example normal, warning, or error.

**Storage Virtualizer Volume Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Virtualizer Volume Native Status**

Shows the level of access of nodes in the system to the storage virtualizer managed disk or volume. The level of access can be online, offline, degraded, or excluded.

**Performance metrics for storage virtualizer volumes:**

You can include performance data for volumes, caches, and other data for storage virtualizer volumes in performance reports.

## Volume data

You can create performance reports that include the following information:

### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### **Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

### **Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

### **Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

## **Volume Utilization**

The average percentage of time that the volume is busy.

### **Maximum Read I/O Rate (ops/s)**

The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

### **Maximum Write I/O Rate (ops/s)**

The maximum write I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

### **Maximum Read Response Time (ms/op)**

The maximum read response time of the volume in milliseconds per

operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Write Response Time (ms/op)**

The maximum write response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Write I/O Size (KiB)**

The maximum size of write I/O operations in kibibytes. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Volume Utilization Percentage**

The maximum period in percentage that the volume is busy. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Maximum Read Cache Hit Percentage**

The maximum number of cache hits in percentage for sequential and nonsequential read operations for the volume over a specified interval. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Write Cache Delay I/O Rate (ops/s)**

The maximum number of delayed I/O operations per second for the volume, where the delay occurred because of space constraints or other conditions.

This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.



**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

**Nonpreferred Node Usage Percentage**

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

**Performance metrics for volumes on storage systems that are configured to virtualize storage:**

A storage environment can include a Hitachi Universal Storage Platform V system that is configured to virtualize storage. You can include performance data for volumes and caches for storage system volumes on these systems in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

## Data for storage virtualizer nodes in performance reports

You can include properties, and other information about storage virtualizer nodes in performance reports.

### Information about storage virtualizer nodes

You can create performance reports that include the following information:

#### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### Storage Virtualizer Node Name

The name that was assigned to the storage virtualizer node when it was added to the system.

#### Storage Virtualizer I/O Group Name

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

### Component properties

You can create performance reports that include the following information:

#### Storage Virtualizer Node IP Address

The IP address of the resource.

#### Storage Virtualizer Node WWN

The worldwide name of the node.

#### Storage Virtualizer Node Configuration

The number of the node for the resource. This value can be either 0 or 1.

### Status information

You can create performance reports that include the following information:

#### Storage Virtualizer Node Status

Shows whether a storage virtualizer node is online.

### Performance metrics for storage virtualizer nodes:

You can include port data, front end data, cache data, back-end data, and other data for storage virtualizer nodes in performance reports.

#### Port data

You can create performance reports that include the following information:

#### Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port.

#### Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data.

#### Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

**Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

**Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

**Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

**Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

**Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

**Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

**Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from

another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Local Node Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

**Overall Port to Local Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Remote Node Response Time (ms/op)**

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

**Overall Port to Remote Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

## Front end data

You can create performance reports that include the following information:

### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### **Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

### **Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

### **Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

## Cache data

You can create performance reports that include the following information:

### **Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### **Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.



## Back-end data

You can create performance reports that include the following information:

### Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

### Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

### Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

### Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

### Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

### Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

### Overall Back-End Queue Time (ms/op)

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

**Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

## **Data for storage virtualizer I/O groups in performance reports**

You can include general information, capacity data, and properties of storage virtualizer I/O groups in performance reports.

### **Information about storage virtualizer I/O groups**

You can create performance reports that include the following information:

#### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### **Storage Virtualizer I/O Group Name**

The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

### **Capacity and usage data**

You can create performance reports that include the following information:

#### **Storage Virtualizer I/O Group Mirroring Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide volume mirroring functions.

#### **Storage Virtualizer I/O Group Mirroring Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide volume mirroring functions.

#### **Storage Virtualizer I/O Group FlashCopy Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide FlashCopy mirroring functions.

#### **Storage Virtualizer I/O Group FlashCopy Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide FlashCopy mirroring functions.

#### **Storage Virtualizer I/O Group Remote Copy Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide mirroring functions by using a remote copy.

#### **Storage Virtualizer I/O Group Remote Copy Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide mirroring functions by using a remote copy.

#### **Storage Virtualizer I/O Group RAID Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide RAID mirroring functions.

#### **Storage Virtualizer I/O Group RAID Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide RAID mirroring functions.

### **Component properties**

You can create performance reports that include the following information:

**Storage Virtualizer I/O Group Maintenance**

Shows whether the I/O group is in maintenance mode. If this value is Yes, the I/O group is in maintenance mode.

**Storage Virtualizer I/O Group Is Compression Active**

Shows whether the compression feature is enabled on the volumes in the storage virtualizer I/O group. If this value is Yes, the compression feature is enabled.

**Storage Virtualizer I/O Group Is Compression Supported**

Shows whether the compression feature is available for volumes in the storage virtualizer I/O group. If this value is Yes, the compression feature is available.

**Performance metrics for storage virtualizer I/O groups:**

You can include performance data for ports, volumes, caches, back-end arrays, and other data for storage virtualizer I/O groups in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

**Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

**Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

**Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

**Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

**Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

**Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

**Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Local Node Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

**Overall Port to Local Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Remote Node Response Time (ms/op)**

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

**Overall Port to Remote Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.



**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

## Other data

You can create performance reports that include the following information:

### **Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

### **Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

### **Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

### **Peak Read Response Time (ms)**

The longest time to complete a read operation.

### **Peak Write Response Time (ms)**

The longest time to complete a write operation.

### **Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

### **Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

### **Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

### **Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases,

this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

**Nonpreferred Node Usage Percentage**

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

**Data for storage virtualizer ports in performance reports**

You can include properties and other information about storage virtualizer ports in performance reports.

**Information about storage virtualizer ports**

You can create performance reports that include the following information:

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer Node Name**

The name that was assigned to the storage virtualizer node when it was added to the system.

**Storage Virtualizer Port Name**

An alphanumeric string that identifies a storage virtualizer port.

**Component properties**

You can create performance reports that include the following information:

**Storage Virtualizer Port Number**

The port number on the resource.

**Storage Virtualizer Port Type**

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N\_Port, F\_Port, or another type of port.

**Storage Virtualizer Port Speed (GiB)**

The speed of a port, which is measured in GiB per second.

**Status information**

You can create performance reports that include the following information:

**Storage Virtualizer Port Status**

The condition of the resource, for example normal, warning, or error.

**Storage Virtualizer Port Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Virtualizer Port Enabled State**

Shows whether a port is enabled, disabled, or is enabled but offline.

**Performance metrics for storage virtualizer ports:**

You can include performance data about storage virtualizer ports and port errors in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

**Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

**Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

**Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

**Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

**Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

**Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

**Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.



## Port error data

You can create performance reports that include the following information:

### Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

### Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

### Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

### CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

### Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

### Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

### Zero Buffer-to-Buffer Credit Timer (microseconds)

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

### Zero Buffer-to-Buffer Credit Percentage

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

## **Performance metrics for ports on storage systems that are configured to virtualize storage:**

A storage environment can include a Hitachi Universal Storage Platform V system that is configured to virtualize storage. You can include performance data for ports on these systems in performance reports.

### **Port data**

You can create performance reports that include the following information:

#### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

#### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

#### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

#### **Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

#### **Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

#### **Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

#### **Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

#### **Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

#### **Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

## **Data for storage virtualizer pools in performance reports**

You can include general information, capacity data, properties, and other information about storage virtualizer pools in performance reports.

### **Information about storage virtualizer pools**

You can create performance reports that include the following information:

#### **Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### **Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

## Capacity and usage data

You can create performance reports that include the following information:

### **Storage Virtualizer Pool Capacity (GiB)**

The total amount of storage space in a pool.

### **Storage Virtualizer Pool Available Space (GiB)**

The amount of space in a pool that is not reserved for volumes.

### **Storage Virtualizer Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

### **Storage Virtualizer Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

### **Storage Virtualizer Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

### **Storage Virtualizer Pool Extent Size (MiB)**

The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

### **Storage Virtualizer Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

### **Storage Virtualizer Pool Number of Storage Volumes**

The number of volumes in the storage virtualizer pool.

### **Storage Virtualizer Pool Number of MDisks**

The number of managed disks on a resource.

### **Storage Virtualizer Pool Physical Allocation Percentage**

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

### **Storage Virtualizer Pool Virtual Allocation Percentage**

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

### **Storage Virtualizer Pool Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{total volume capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50

GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool Unallocated Volume Space (GiB)**

The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

total volume capacity - allocated space

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

**Storage Virtualizer Pool Unused Volume Space (GiB)**

The amount of space that is allocated to the volumes in a pool and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

allocated space - used pool space

**Storage Virtualizer Pool Unallocatable Volume Space (GiB)**

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

**Storage Virtualizer Pool Used Space Percentage**

The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

**Storage Virtualizer Pool Compressed Virtual Capacity (GiB)**

The total virtual capacity of all the volumes that are compressed in a storage virtualizer pool.

**Storage Virtualizer Pool Compressed Capacity (GiB)**

The amount of storage space that is used by compressed volumes in a storage virtualizer pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

**Storage Virtualizer Pool Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

## Component properties

You can create performance reports that include the following information:

### Storage Virtualizer Pool Is Encryptable

Shows whether the storage virtualizer pool can be encrypted. If this value is Yes, the storage virtualizer pool can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer Pool Is Encrypted

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

### Storage Virtualizer Pool Is Solid State

Shows whether there are solid-state drives in the pool. This property can contain the following values:

**Mixed** The pool contains both hard disk drives and solid-state drives.

**Non solid state**

The pool contains no solid-state drives.

**Solid state**

The pool contains a least one solid-state drive.

### Storage Virtualizer Pool Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

### Storage Virtualizer Pool RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

### Storage Virtualizer Pool Warning Level

The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### Storage Virtualizer Pool Is Compression Active

Shows whether the compression feature is enabled on the storage virtualizer pool. If this value is Yes, the compression feature is enabled.

## Status information

You can create performance reports that include the following information:

### Storage Virtualizer Pool Status

The condition of the resource, for example normal, warning, or error.

### Storage Virtualizer Pool Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

## Performance metrics for storage virtualizer pools:

You can include performance data for volumes and back-end arrays for storage virtualizer pools in performance reports.

## Volume data

You can create performance reports that include the following information:

### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### **Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

### **Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

### **Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

### **Maximum Read I/O Rate (ops/s)**

The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

## Back-end array data

You can create performance reports that include the following information:

### **Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.



**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Maximum Back-End Read I/O Rate (ops/s)**

The maximum number of read operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Back-End Write I/O Rate (ops/s)**

The maximum number of write operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Back-End Read Response Time (ms/op)**

The maximum number of milliseconds for the back-end storage resources to respond to a read operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Back-End Write Response Time (ms/op)**

The maximum number of milliseconds for the back-end storage resources to respond to a write operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

---

## Predefined reports about resource relationships

You can run predefined reports about the end-to-end relationship from a server to back-end storage. You can also run reports about the end-to-end relationship from file systems on servers to volumes on storage pools.

### Running the Storage Resource Relationships Summary report

To show the relationships between storage resources, run the Storage Resource Relationships Summary report.

To run the report with the default settings, click **Finish**.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Resource Relationships Summary**.
3. Optional: Select one or more of the following:

- a. Servers
  - b. Hypervisors
  - c. NAS resources
  - d. Storage virtualizers
  - e. Storage systems
4. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

**Restriction:** To create a report when you select a single resource, the selected resource must be connected to other resources.

**Related reference:**

“Storage Resource Relationships Summary report”

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

## Storage Resource Relationships Summary report

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

### Report output

Information is provided about related resources and the devices, such as managed disks (MDisks), that are related to resources.

#### Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

#### Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

#### NAS Name

The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

#### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

#### Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

#### Storage Virtualizer Volume Assigned Host Connection

The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

#### Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

**Storage Virtualizer MDisk Name**

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Volume Assigned Host Connection**

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Volume	Performance of One Volume
Managed disk	Performance of One Managed Disk
Storage pool	Most Active Pools, Performance of One Pool

**Related tasks:**

“Running the Storage Resource Relationships Summary report” on page 627

To show the relationships between storage resources, run the Storage Resource Relationships Summary report.

**Running the Storage Resource Relationships Summary (Configurable) report**

To show the relationships between resources, run the Storage Resource Relationships Summary (Configurable) report. You can show or hide resources, and you can hide resources that are not related to other resources or storage devices.

To run the report with the default settings, click **Finish**. Servers, storage systems, storage virtualizers, and their related resources are shown in the report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

**Tip:** To create a report for a single resource, the selected resource must be connected to other storage devices or resources. To create a report about two or more selected resources, the resources that you select must be related.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Resource Relationships Summary (Configurable)**.
3. Optional: Select the storage resources that you want to show in the report.
4. Optional: Specify whether you want to show or hide a resource that is related to other storage devices such as a file system, but not to other resources. For example, a server has a related file system. You can show the server and its related storage devices in the report, or you can hide the resource and its related storage devices in the report.
5. Optional: Click **Finish** to run the report.
6. Optional: Click **Next**. In the Select storage resources page, select one or more of the following:
  - a. Servers
  - b. Hypervisors
  - c. NAS resources
  - d. Storage virtualizers
  - e. Storage systems
7. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

**Related reference:**

“Storage Resource Relationships Summary (Configurable) report”

Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

## Storage Resource Relationships Summary (Configurable) report

Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

### Report output

Information is provided about related resources and the devices, such as managed disks (MDisks), that are related to resources.

**Server Name**

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

**Server Disk Name**

The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

**Logical Volume Path**

The path to a logical volume on a resource, for example `/dev/hd1`.

**File System Mount Point**

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

**Hypervisor Name**

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

**Hypervisor Disk Name**

The path that the operating system uses for the server disk.

**NAS Name**

The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

**NAS Export Name**

The name of the exported file system.

**Storage Virtualizer Name**

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

**Storage Virtualizer Volume Assigned Host Connection**

The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

**Storage Virtualizer Volume Name**

The name that was assigned to the storage virtualizer volume when it was added to the system.

**Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Virtualizer MDisk Name**

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

**Storage Volume Assigned Host Connection**

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Volume	Performance of One Volume
Managed disk	Performance of One Managed Disk
Storage pool	Most Active Pools, Performance of One Pool

### Related tasks:

“Running the Storage Resource Relationships Summary (Configurable) report” on page 629

To show the relationships between resources, run the Storage Resource Relationships Summary (Configurable) report. You can show or hide resources, and you can hide resources that are not related to other resources or storage devices.

## Running the File System to Volume Relationships report

To show the end-to-end relationship from file systems on servers to volumes on storage pools, run the File System to Volume Relationships.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **File System to Volume Relationships**.
3. Select one or more servers.
4. Select the storage virtualizers and storage systems that the storage pools are in. You can select both storage virtualizers and storage systems, storage virtualizers only, or storage systems only.
5. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

**Restriction:** To create a report when you select a single resource, the selected resource must be connected to other resources.

### Related reference:

“File System to Volume Relationships report”

Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

## File System to Volume Relationships report

Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

## Report output

The information that is provided about the servers, storage virtualizers, storage systems, pools, and volumes is as follows:

### Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

### File System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

### Server Disk Name

The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

### File System Capacity (GiB)

The amount of storage space on the file system of the resource.

### File System Available Space (GiB)

The amount of unused storage space in the file system of the resource.

### File System Available Space Percentage

The percentage of space in the file system that is unused.

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

### Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

### Storage Virtualizer Volume Is Thin Provisioned

Shows whether a volume or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

### Storage Virtualizer Volume Allocated Space Percentage

The percentage of space on a volume that is reserved.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

### Storage Virtualizer Volume Used Space Percentage

The percentage of reserved space that is being used on a volume.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

### Storage Virtualizer Pool Name

The name that was assigned to the pool when the pool was added to the system.

### Storage Virtualizer Pool Available Space (GiB)

The amount of space in a pool that is not reserved for volumes.

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Storage Volume Is Thin Provisioned**

Shows whether a volume or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Volume Allocated Space Percentage**

The percentage of space on a volume that is reserved.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

**Storage Volume Used Space Percentage**

The percentage of reserved space that is being used on a volume.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Pool Available Space (GiB)**

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$\text{pool capacity} - \text{used space}$

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

**Related tasks:**

“Running the File System to Volume Relationships report” on page 632

To show the end-to-end relationship from file systems on servers to volumes on storage pools, run the File System to Volume Relationships.

---

## Predefined reports about switches and switch ports

You can run predefined reports to analyze and compare the performance of switches and switch ports.

### Predefined reports about switches

You can run predefined reports to analyze and compare the performance of switches, and to see information about the ports that are available from switches.

**Running the Availability of Switch Ports report**

To see information about the ports and connected ports that are available from switches, run the Availability of Switch Ports report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.



3. Click **Switches**.
4. Click **Availability of Switch Ports**.
5. Select the switches, and then select a sort order.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Availability of Switch Ports report”

Shows a chart of the 20 switches with the most ports and connected ports, and a table with the properties of the switches.

### **Availability of Switch Ports report**

Shows a chart of the 20 switches with the most ports and connected ports, and a table with the properties of the switches.

### **Charts**

The chart shows the number of ports for each switch that are not online, and the number of ports that are online and connected to other ports.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

For each switch, the following information is provided:

**Switch Name**

The logical name of the switch where a Fiber Channel port is located, or the name that was defined when the switch was added for monitoring. If neither name is available, Tivoli Storage Productivity Center uses the WWN of the switch.

**Switch Location**

The physical location of the switch. The location is defined when a switch is added to Tivoli Storage Productivity Center. In the web-based GUI, you can add or edit the location of the switch on the **General** tab of the properties notebook.

**Switch WWN**

The World Wide Name (WWN) of the switch. A WWN is the unique 64-bit identifier for the switch.

**Switch Is Virtual**

Shows whether a switch is a logical switch. If this value is Yes, the switch is a logical switch.

**Switch IP Address**

The IP address of the resource.

**Switch Vendor**

The vendor who supplied the resource.

**Switch Model**

The model name or model number of the resource.

**Switch Ports**

The total number of ports on the switch that are not online, and ports that are online and connected to other ports.

**Switch Connected Ports**

The number of ports that are connected to a storage resource, where the storage resource can be a storage system, server, or hypervisor. This value is only available when a switch is viewed as a related resource of a storage system, server, or hypervisor.

**Switch Last Data Collection**

The date and time when storage statistics were last collected from the resource.

**Switch Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

**Related tasks:**

“Running the Availability of Switch Ports report” on page 634

To see information about the ports and connected ports that are available from switches, run the Availability of Switch Ports report.

**Running the Performance of One Switch report**

To see information about the performance of a switch, run the Performance of One Switch report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switches**.
4. Click **Performance of One Switch**.
5. Select a switch.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Switch report”

Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

**Performance of One Switch report**

Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

## Charts

The charts show the following performance metrics for the ports on the switch:

- The send frame rates, the receive frame rates, and the total frame rates
- The send data rates, the receive data rates, and the total data rates
- The send bandwidth percentages, the receive bandwidth percentages, and the overall bandwidth percentages

## Report output

The following information is provided in the report for the period that you specified:

### Switch Name

The logical name of the switch or the name that was defined when the switch was added to Tivoli Storage Productivity Center. If neither name is available, the WWN of the switch is displayed.

### Switch Location

The physical location of the switch. The location is defined when a switch is added to Tivoli Storage Productivity Center. In the web-based GUI, you can add or edit the location of the switch on the **General** tab of the properties notebook.

### Port Send Frame Rate (frames/s)

The average number of frames per second that are sent through the port.

### Port Receive Frame Rate (frames/s)

The average number of frames per second that are received by the port.

### Total Port Frame Rate (frames/s)

The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

### Port Send Data Rate (MiB/s)

The average rate at which data is sent through the port. The rate is measured in MiB per second.

### Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second.

### Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

### Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

### Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

### Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

## Related reports

To open the Most Active Switch Ports report, click the name of the switch in the report table.

### Related tasks:

“Running the Performance of One Switch report” on page 636

To see information about the performance of a switch, run the Performance of One Switch report.

## Running the Compare Performance of Multiple Switches report

Use the Compare Performance of Multiple Switches report to compare up to four performance metrics for multiple switches. For example, you can compare bandwidth percentages, data rates, and frame rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switches**.
4. Click **Compare Performance of Multiple Switches**.
5. Select the switches, and then click **Next**.
6. Select the category of performance metric that you want to include in the report, and then select the performance metric.
7. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

### Related reference:

“Compare Performance of Multiple Switches report”

Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

## Compare Performance of Multiple Switches report

Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

## Charts

The charts show the performance metrics for the switches that you selected for the period that you specified.

## Report output

The following information is provided in the report:

### Performance metrics

Shows the data for each performance metric.

### Switch names

Shows the names of the switches, and the performance metric data for each switch.

### Related tasks:

“Running the Compare Performance of Multiple Switches report” on page 638  
Use the Compare Performance of Multiple Switches report to compare up to four performance metrics for multiple switches. For example, you can compare bandwidth percentages, data rates, and frame rates.

## Running the Compare Performance of One Switch over Time Ranges report

To compare a performance metric for one switch in two date ranges, run the Compare Performance of One Switch over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switches**.
4. Click **Compare Performance of One Switch over Time Ranges**.
5. Select the switch.
6. Select the category of performance metric about which you want a report, and then select the performance metric.
7. Select an interval.
8. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

**Tip:** To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.

9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Compare Performance of One Switch over Time Ranges report”

Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

## Compare Performance of One Switch over Time Ranges report

Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

## Charts

The chart shows the performance data of the switch on the performance metric you selected, over the time periods that you selected.

## Report output

The following information is provided in the report:

### Switch Name

The logical name of the switch or the name that was defined when the switch was added to Tivoli Storage Productivity Center. If neither name is available, the WWN of the switch is displayed.

### Day in Range

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

### Hour in Range

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

### Minute in Range

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

### First Date Range

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

### Second Date Range

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

### Related tasks:

“Running the Compare Performance of One Switch over Time Ranges report” on page 639

To compare a performance metric for one switch in two date ranges, run the Compare Performance of One Switch over Time Ranges report.

## Predefined reports about switch ports

You can run predefined reports to analyze and compare the performance of switch ports.

## Running the Most Active Switch Ports report

To see which ports are most active on switches, run the Most Active Switch Ports report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Most Active Switch Ports**.
5. Select the switches, and then select a sort order.
6. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Most Active Switch Ports report”

Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

## Most Active Switch Ports report

Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

### Charts

The bar chart shows the most active ports on the switch. By default, ports are sorted by the total frame rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

The information that is provided about the most active ports on switches is as follows:

#### Switch Name

The logical name of the switch or the name that was defined when the switch was added to Tivoli Storage Productivity Center. If neither name is available, the WWN of the switch is displayed.

#### Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

**Switch Port Number**

The number of the port on the switch.

**Switch Port WWPN**

The worldwide port name of the port on the switch.

**Total Port Frame Rate (frames/s)**

The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Switch	Performance of One Switch
Switch port	Performance of One Switch Port

**Related tasks:**

“Running the Most Active Switch Ports report” on page 641

To see which ports are most active on switches, run the Most Active Switch Ports report.

**Running the Performance of One Switch Port report**

To see information about the performance of a switch port, run the Performance of One Switch Port report.

You must select a switch and a switch port. If the port is on a blade, you must select the number of the slot on the switch to which the blade is attached.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Performance of One Switch Port**.
5. Select a switch.
6. Optional: Select the number of the blade slot. If the port is on a blade, you must select the number of the blade slot.
7. Select a port.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.



**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Switch Port report”

Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

### **Performance of One Switch Port report**

Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

### **Charts**

The charts show the following performance statistics for the port:

- The send frame rate, the receive frame rate, and the total frame rate
- The send data rate, the receive data rate, and the total data rate
- The send bandwidth percentage, the receive bandwidth percentage, and the overall bandwidth percentage
- The error rates

### **Report output**

The following information is provided in the report for the period that you specified:

**Switch Name**

The logical name of the switch or the name that was defined when the switch was added to Tivoli Storage Productivity Center. If neither name is available, the WWN of the switch is displayed.

**Switch Blade Slot Number**

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

**Switch Port Number**

The number of the port on the switch.

**Switch Port WWPN**

The worldwide port name of the port on the switch.

**Port Send Frame Rate (frames/s)**

The average number of frames per second that are sent through the port.

**Port Receive Frame Rate (frames/s)**

The average number of frames per second that are received by the port.

**Total Port Frame Rate (frames/s)**

The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Error Frame Rate (count/s)**

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

**Discarded Frame Rate (count/s)**

The average number of frames per second that are discarded because host buffers are unavailable for the port.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Discarded Class 3 Frame Rate (count/s)**

The average number of class 3 frames per second that are discarded.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Related reports**

To open the Most Active Switch Ports report, click the name of the switch in the report table.

**Related tasks:**

“Running the Performance of One Switch Port report” on page 642

To see information about the performance of a switch port, run the Performance of One Switch Port report.

## Running the Compare Performance of Multiple Switch Ports report

Use the Compare Performance of Multiple Switch Ports report to compare up to four performance metrics for multiple switch ports. For example, you can compare bandwidth percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Compare Performance of Multiple Switch Ports**.
5. Select the switches.
6. Select switch ports. To select switch ports, complete the following steps:
  - a. Type a keyword. For example, type `example1` to find the names of switch ports that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find switch ports with names begin with 'a' or 'A'. Type `%a` to find switch ports with names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the switch ports on a resource.
  - b. Click **Search**. If the search is successful, the switch ports are displayed in the **Results** list.
  - c. Select the switch ports that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

12. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

### Related reference:

“Compare Performance of Multiple Switch Ports report”

Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

## Compare Performance of Multiple Switch Ports report

Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

## Charts

The charts show the performance metrics for the switch ports that you selected for the period that you specified.

## Report output

The following information is provided in the report:

### Performance metrics

Shows the data for each performance metric.

### Switch names

Shows the names of the switches, and the performance metric data for each switch.

### Slot numbers

Shows the slot numbers, and the performance metric data for each slot.

### Switch port numbers

Shows the port numbers, and the performance metric data for each switch port number.

### Related tasks:

“Running the Compare Performance of Multiple Switch Ports report” on page 645  
Use the Compare Performance of Multiple Switch Ports report to compare up to four performance metrics for multiple switch ports. For example, you can compare bandwidth percentages, data rates, and I/O rates.

## Running the Compare Performance of One Switch Port over Time Ranges report

To compare a performance metric for one switch port in two date ranges, run the Compare Performance of One Switch Port over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Compare Performance of One Switch Port over Time Ranges**.
5. Select the switch that the port is on.
6. If the switch is on a blade, select the slot number of the blade that the port is on.
7. Select the number of the port.
8. Select the category of performance metric about which you want a report, and then select the performance metric.
9. Select an interval.
10. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

**Tip:** To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.

11. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Compare Performance of One Switch Port over Time Ranges report”  
Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

**Compare Performance of One Switch Port over Time Ranges report**

Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

**Charts**

The chart shows the performance data of the switch port on the performance metric you selected, over the time periods that you selected.

**Report output**

The following information is provided in the report:

**Switch Name**

The logical name of the switch or the name that was defined when the switch was added to Tivoli Storage Productivity Center. If neither name is available, the WWN of the switch is displayed.

**Switch Blade Slot Number**

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

**Switch Port Number**

The number of the port on the switch.

**Switch Port WWPN**

The worldwide port name of the port on the switch.

**Day in Range**

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

**Hour in Range**

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

**Minute in Range**

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

**First Date Range**

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was

collected. The column with the performance metric name shows the data for the performance metric for the first date range.

### **Second Date Range**

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

### **Related tasks:**

“Running the Compare Performance of One Switch Port over Time Ranges report” on page 646

To compare a performance metric for one switch port in two date ranges, run the Compare Performance of One Switch Port over Time Ranges report.

---

## **Predefined reports about groups**

You can run predefined reports to analyze the capacity of groups. A group is a set of logically related volumes, file systems, and shares. For example, a group that represents a business critical application might include the volumes, file systems, and shares that provide storage to the application.

### **Running the Capacity of One Group report**

To see information about the block-level and file-level capacity and space of a group, run the Capacity of One Group report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Groups**.
3. Click **Capacity of One Group**.
4. Select a group.
5. Optional: Select a sort order.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### **Related reference:**

“Capacity of One Group report”

Shows charts and tables of the capacity, available space, and used space statistics for each type of resource in a group. The report shows separate charts and tables for block-level and file-level resources.

### **Capacity of One Group report**

Shows charts and tables of the capacity, available space, and used space statistics for each type of resource in a group. The report shows separate charts and tables for block-level and file-level resources.

### **Charts**

The charts show the following block-level and file-level capacity statistics for a group:

- Total capacity

- Total available space
- Total used space

If there are any block-level storage devices in the group, the chart also shows the total allocated space for those devices.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

## Report output

The report provides the following information about the group:

### Group Name

The name that was assigned to the group by the user when the group was created.

### Resource Name

For a server or NAS file server, this property shows the fully qualified domain name of the server or NAS file server.

For a block-level storage resource, this property shows the user-defined name of the resource.

### Resource Type

For a server or NAS file server, this property shows the operating system that is running on the server or NAS file server.

For a block-level storage resource, this property shows the type of resource. For example, the property can be **IBM System Storage DS8800**, **Engenio DS4000**, or another type of resource.

### Total File Capacity (GiB)

The total amount of file-level storage space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

### Total File Used Space (GiB)

The total amount of file-level storage space that is used by the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

### Total File Available Space (GiB)

The total amount of file-level unused space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

### Total Block Capacity (GiB)

The total amount of block-level storage space that is committed to storage system volumes or storage virtualizer volumes in a group.

### Total Block Used Space (GiB)

The total amount of block-level allocated space that is used by storage system volumes or storage virtualizer volumes in a group.

### Total Block Available Space (GiB)

The total amount of block-level unused space that is available on storage system volumes or storage virtualizer volumes in a group.

### Total Block Allocated Space (GiB)

The total amount of block-level space that is reserved for the storage system volumes or storage virtualizer volumes in a group. The allocated space includes both thin-provisioned and standard volumes.

**Related tasks:**

“Running the Capacity of One Group report” on page 648

To see information about the block-level and file-level capacity and space of a group, run the Capacity of One Group report.

## Running the Groups Capacity report

To see information about the capacity and space of groups, run the Groups Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Groups**.
3. Click **Groups Capacity**.
4. Optional: Select one or more groups.
5. Optional: Select a sort order.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Groups Capacity report”

Shows a chart of the 20 groups with the most capacity, available space, or used space, and a table with details of all groups. Use the report to review the space allocation in your groups.

## Groups Capacity report

Shows a chart of the 20 groups with the most capacity, available space, or used space, and a table with details of all groups. Use the report to review the space allocation in your groups.

### Charts

The bar chart shows the following block-level and file-level capacity statistics for groups:

- Total available space
- Total used space

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

For each group, the following information is provided:

**Group Custom Tag 1**

User-defined text that is associated with a group.

**Group Name**

The name that was assigned to the group by the user when the group was created.



**Total Capacity (GiB)**

The total amount of block-level and file-level storage space that is committed to volumes, file systems, and shares.

**Total Used Space (GiB)**

The total amount of block-level and file-level space that is used by volumes, file systems, and shares.

**Total Available Space (GiB)**

The total amount of block-level and file-level space that is available on volumes, file systems, and shares.

**Total Block Capacity (GiB)**

The total amount of block-level storage space that is committed to storage system volumes or storage virtualizer volumes in a group.

**Total Block Used Space (GiB)**

The total amount of block-level allocated space that is used by storage system volumes or storage virtualizer volumes in a group.

**Total Block Available Space (GiB)**

The total amount of block-level unused space that is available on storage system volumes or storage virtualizer volumes in a group.

**Total Block Allocated Space (GiB)**

The total amount of block-level space that is reserved for the storage system volumes or storage virtualizer volumes in a group. The allocated space includes both thin-provisioned and standard volumes.

**Total File Capacity (GiB)**

The total amount of file-level storage space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

**Total File Used Space (GiB)**

The total amount of file-level storage space that is used by the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

**Total File Available Space (GiB)**

The total amount of file-level unused space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

**Related tasks:**

“Running the Groups Capacity report” on page 650

To see information about the capacity and space of groups, run the Groups Capacity report.

---

## Predefined reports about hypervisors

You can run predefined reports to analyze the capacity and performance of hypervisors.

### Running the Hypervisors Capacity report

To see information about the capacity of disks on hypervisors, run the Hypervisors Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Hypervisors Capacity**.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a storage resources sort order.
7. Optional: Select a sort order.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Hypervisors Capacity report”

Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

## Hypervisors Capacity report

Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

### Charts

The bar chart shows the used space and available space on hypervisors.

### Report output

For each hypervisor, the following information is provided:

**Hypervisor Name**

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

**Hypervisor Location**

The physical location of a hypervisor. The location is defined when a hypervisor is added to Tivoli Storage Productivity Center. You can add or edit the location of the hypervisor in the **Properties** pane of the hypervisor in the web-based GUI.

**Hypervisor Cluster Name**

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

**Hypervisor Total Disk Capacity (GiB)**

The amount of storage space that is on a hypervisor.

**Hypervisor Total Disk Available Space (GiB)**

The amount of unused storage space on all disks on a hypervisor.

**Hypervisor Total Disk Used Space (GiB)**

The amount of storage space that is used on all disks on a hypervisor.

**Hypervisor Total Disk Used Space (%)**

The amount of storage space that is used on a hypervisor.

**Hypervisor Data Store Capacity (GiB)**

The amount of storage space that is on the hypervisor data store.

**Hypervisor Data Store Available Space (GiB)**

The amount of unused storage space on a hypervisor data store.

**Hypervisor Data Store Used Space (GiB)**

The amount of storage space that is used on a hypervisor data store.

**Hypervisor Data Store Used Space (%)**

The percentage of storage space that is used on all disks on a hypervisor.

**Related reports**

To open related reports, click the name of the hypervisor and then click the name of the report. You can open the following related reports:

- Hypervisor Disks Capacity
- Hypervisor Data Stores Capacity

**Related tasks:**

“Running the Hypervisors Capacity report” on page 651

To see information about the capacity of disks on hypervisors, run the Hypervisors Capacity report.

**Running the Most Active Hypervisors report**

To see which hypervisors are the most active, run the Most Active Hypervisors report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Most Active Hypervisors**.
4. Select the hypervisors, and then select a sort order.
5. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Hypervisors report” on page 654

Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

## Most Active Hypervisors report

Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

### Charts

The bar chart shows the most active hypervisors. By default, hypervisors are sorted by the total I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

The information that is provided about the most active hypervisors is as follows:

#### Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

#### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

#### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

#### Volume Utilization

The average percentage of time that the volume is busy.

### Related reports

To open the Summarized Performance of Volumes by Hypervisor report, click the name of the hypervisor in the report table.

#### Related tasks:

“Running the Most Active Hypervisors report” on page 653

To see which hypervisors are the most active, run the Most Active Hypervisors report.

## Running the Summarized Performance of Volumes by Hypervisor report

To see the load that a particular hypervisor adds to your storage environment, use the Summarized Performance of Volumes by Hypervisor report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.

2. Click **Hypervisors**.
3. Click **Summarized Performance of Volumes by Hypervisor**.
4. Select one or more hypervisors.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Summarized Performance of Volumes by Hypervisor report”

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

## Summarized Performance of Volumes by Hypervisor report

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

### Charts

The charts show the following performance metrics for the volumes on a hypervisor:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write data rates, and the total data rates
- The read cache percentages and the write cache percentages

### Report output

The following information is provided in the report for the period that you specified:

**Hypervisor Name**

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

**Day, hour, or time**

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation, across all volumes on the hypervisor. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second, across all volumes on the hypervisor. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the hypervisor.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the hypervisor. The average value for each volume is added to give a total data rate for all volumes on the hypervisor.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation, across all volumes on the hypervisor.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation, across all volumes on the hypervisor.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations, across all volumes on the hypervisor. The average value for each volume is added to give a read data rate for all volumes on the hypervisor.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations, across all volumes on the hypervisor. The average value for each volume is added to give a write data rate for all volumes on the hypervisor.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second, which is totaled for all volumes on the hypervisor. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second, which is totaled for all volumes on the hypervisor. This value includes both sequential and nonsequential write operations.

**Overall Read Cache Hit Percentage**

The average percentage of all read operations that find data in the cache, across all volumes on the hypervisor. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The average percentage of all write operations that are handled in the cache, across all volumes on the hypervisor. This value includes both sequential and nonsequential write operations.

**Related reports**

To open the Performance of Volumes by Hypervisor report, click the name of the hypervisor in the report table.

**Related tasks:**

“Running the Summarized Performance of Volumes by Hypervisor report” on page 654

To see the load that a particular hypervisor adds to your storage environment, use the Summarized Performance of Volumes by Hypervisor report.

## Running the Performance of Volumes by Hypervisor report

To see performance metrics for the volumes on a hypervisor, use the Performance of Volumes by Hypervisor report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Performance of Volumes by Hypervisor**.
4. Select one or more hypervisors.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Performance of Volumes by Hypervisor report”

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

## Performance of Volumes by Hypervisor report

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

### Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. The chart shows the contribution of each volume to the overall load on the hypervisor.
- The total data rates. The chart shows the contribution of each volume to the overall load on the hypervisor.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

## Report output

The following information is provided in the report for the period that you specified:

### Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage Pool Name

The name that was assigned to the pool when it was added to the system.

### Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

### Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Volume Utilization

The average percentage of time that the volume is busy.

### Related tasks:

“Running the Performance of Volumes by Hypervisor report” on page 657  
To see performance metrics for the volumes on a hypervisor, use the Performance of Volumes by Hypervisor report.

## Running the Hypervisor Data Stores Capacity report

To see information about the capacity of data stores on hypervisors, run the Hypervisor Data Stores Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Data Stores Capacity**.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.



6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Hypervisor Data Stores Capacity report”

Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

## Hypervisor Data Stores Capacity report

Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

### Charts

The bar chart shows the used space and available space for data stores on hypervisors.

### Report output

For each hypervisor, the following information is provided:

**Hypervisor Cluster Name**

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

**Hypervisor Data Store**

The path to the VMware ESX data store on the hypervisor, for example `datastore_svc3c`.

**Hypervisor Data Store Cluster Name**

The name of the cluster, if the data store is a member of a data store cluster. If the data store is not a member of a data store cluster, no information is displayed.

**Hypervisor Data Store Type**

The type of file system that the hypervisor data store uses. For example, the type can be NFS or VMFS.

**Hypervisor System Mount Point**

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

**Hypervisor Data Store Capacity (GiB)**

The amount of storage space that is on the hypervisor data store.

**Hypervisor Data Store Used Space (GiB)**

The amount of storage space that is used on a hypervisor data store.

**Hypervisor Data Store Available Space (GiB)**

The amount of unused storage space on a hypervisor data store.

**Hypervisor Data Store Used Space (%)**

The percentage of storage space that is used on all disks on a hypervisor.

**Hypervisor File System Capacity (GiB)**

The amount of storage space on the file system of the hypervisor.

**Hypervisor File System Available Space (GiB)**

The amount of unused storage space in the file system of the hypervisor.

**Related tasks:**

“Running the Hypervisor Data Stores Capacity report” on page 658

To see information about the capacity of data stores on hypervisors, run the Hypervisor Data Stores Capacity report.

## Running the Hypervisor Disks Capacity report

To see information about the disk capacity on hypervisors, run the Hypervisor Disks Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Hypervisor Disks Capacity**.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Hypervisor Disks Capacity report”

Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

## Hypervisor Disks Capacity report

Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

### Charts

The bar chart shows the used space and available space on hypervisors.

### Report output

For each hypervisor, the following information is provided:

**Hypervisor Cluster Name**

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

**Hypervisor Disk Name**

The path that the operating system uses for the server disk.

**Hypervisor Disk Capacity (GiB)**

The amount of storage space that is on a hypervisor disk.

**Hypervisor Disk Available Space (GiB)**

The amount of unused storage space that is on a hypervisor disk.

**Hypervisor Disk Used Space (GiB)**

The amount of storage space that is used on a hypervisor disk.

**Hypervisor Disk Used Space (%)**

The amount of storage space that is used on a hypervisor disk.

**Hypervisor Disk Status**

The status of a disk on the hypervisor. Use the status to determine the condition of the disk, and if any actions must be taken. For example, if a disk has an Error status, take immediate action to correct the problem. If the disk has an Operational status, then it is operating normally and no further action is required.

**Hypervisor Disk Vendor**

The vendor who supplied the hypervisor disk.

**Hypervisor Disk Model**

The model name or model number of the hypervisor disk.

**Hypervisor Disk Serial Number**

The serial number of the hypervisor disk.

**Hypervisor Disk Firmware Version**

The version number of the firmware that is running on the hypervisor disk.

**Related tasks:**

“Running the Hypervisor Disks Capacity report” on page 660

To see information about the disk capacity on hypervisors, run the Hypervisor Disks Capacity report.

---

## Predefined reports about servers and file systems on servers

You can run predefined reports to analyze the capacity and performance of servers, file systems, and volumes.

### Predefined reports about servers

You can run predefined reports to analyze the capacity and performance of servers and volumes.

#### Running the Servers Capacity report

To see information about the capacity and space of server disks and file systems, run the Servers Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Servers Capacity**.
4. Optional: Select one or more servers.
5. Optional: Select a sort option. You can order the metrics represented in the chart by disk or by file system.

**Tip:** When the report is created, you can change the sort option.

6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Servers Capacity report”

Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

### **Servers Capacity report**

Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

### **Charts**

The bar chart shows the capacity, used space, available space, and percentage of used space for disks on servers. Dots are used to represent the total space that is used by the file system, and the total space that is available to the file system on the server. You can position the cursor over the dot to show the amount of used space or the amount of available space for the file system. The position of the dot that is used to represent the total space that is available to the file system is also used to show the total capacity of the file system.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

For each server, the following information is provided:

**Server Name**

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

**Server Location**

The physical location of a server. The location is defined when a server is added to Tivoli Storage Productivity Center. You can add or edit the location of the server in the **Properties** pane of the server in the web-based GUI.

**Server Total Disk Capacity (GiB)**

The amount of storage space that is on a server.

**Server Total Disk Used Space (GiB)**

The amount of storage space that is used on all disks on a server.

**Server Total Disk Available Space (GiB)**

The amount of unused storage space on all disks on a server.

**Server Total Disk Used Space (%)**

The percentage of used storage space on all disks on the server.

**Server Total File System Capacity (GiB)**

The amount of storage space on the file system of the resource.

**Server Total File System Used Space (GiB)**

The amount of used storage space in the file system of the resource.

**Server Total File System Available Space (GiB)**

The amount of unused storage space in the file system of the resource.

**Server Total File System Used Space (%)**

The amount of used storage space in the file system of the resource.

**Related reports**

To open related reports, click the name of the server and then click the name of the report. You can open the following related reports:

- Server Disks Capacity
- File Systems Capacity

**Related tasks:**

“Running the Servers Capacity report” on page 661

To see information about the capacity and space of server disks and file systems, run the Servers Capacity report.

**Running the Most Active Servers report**

To see which servers are the most active, run the Most Active Servers report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Most Active Servers**.
4. Select the servers, and then select a sort order.
5. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Servers report” on page 664

Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

## Most Active Servers report

Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

### Charts

The bar chart shows the most active servers. By default, servers are sorted by the total I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

The information that is provided about the most active storage systems is as follows:

#### Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

#### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

#### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

#### Volume Utilization

The average percentage of time that the volume is busy.

### Related reports

To open the Summarized Performance of Volumes by Server report, click the name of the server in the report table.

#### Related tasks:

“Running the Most Active Servers report” on page 663

To see which servers are the most active, run the Most Active Servers report.

### Running the Summarized Performance of Volumes by Server report

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Server report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.

3. Click **Summarized Performance of Volumes by Server**.
4. Select one or more servers.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Summarized Performance of Volumes by Server report”

Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

### **Summarized Performance of Volumes by Server report**

Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

### **Charts**

The charts show the following performance metrics for the volumes on a server:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write data rates, and the total data rates
- The read cache percentages and the write cache percentages

### **Report output**

The following information is provided in the report for the period that you specified:

**Server Name**

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

**Day, hour, or time**

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation, across all volumes on the server. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second, across all volumes on the server. This value includes both sequential and

nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the server.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the server. The average value for each volume is added to give a total data rate for all volumes on the server.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation, across all volumes on the server.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation, across all volumes on the server.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations, across all volumes on the server. The average value for each volume is added to give a read data rate for all volumes on the server.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations, across all volumes on the server. The average value for each volume is added to give a write data rate for all volumes on the server.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential write operations.

**Overall Read Cache Hit Percentage**

The average percentage of all read operations that find data in the cache, across all volumes on the server. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

**Related reports**

To open the Performance of Volumes by Server report, click the name of the server in the report table.

**Related tasks:**

“Running the Summarized Performance of Volumes by Server report” on page 664  
To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Server report.



## Running the Performance of Volumes by Server report

To see performance metrics for the volumes on a server, use the Performance of Volumes by Server report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Performance of Volumes by Server**.
4. Select a server.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Performance of Volumes by Server report”

Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

## Performance of Volumes by Server report

Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

## Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. The chart shows the contribution of each volume to the overall load on the server.
- The total data rates. The chart shows the contribution of each volume to the overall load on the server.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

## Report output

The following information is provided in the report for the period that you specified:

### Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Day, hour, or time**

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Related tasks:**

“Running the Performance of Volumes by Server report” on page 667

To see performance metrics for the volumes on a server, use the Performance of Volumes by Server report.

**Running the Server Disks Capacity report**

To see information about the capacity of disks on servers, run the Server Disks Capacity report.

To run the report with the default settings, click **Finish**.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Server Disks Capacity**.
4. Optional: Select one or more servers.
5. Optional: Select a sort order.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Server Disks Capacity report”

Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

### **Server Disks Capacity report**

Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

### **Charts**

The bar chart shows the capacity, used space, available space, and percentage of used space for disks on servers.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

For each server, the following information is provided:

#### **Server Disk Name**

The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

#### **Server Disk Capacity (GiB)**

The amount of storage space that is on a server disk.

#### **Server Disk Available Space (GiB)**

The amount of unused storage space that is on a server disk.

#### **Server Disk Used Space (GiB)**

The amount of storage space that is used on a server disk.

#### **Server Disk Used Space (%)**

The percentage of used storage space on the disk.

#### **Server Disk Status**

The condition of the resource, for example normal, warning, or error.

#### **Server Disk Firmware Version**

The version number of the firmware that is running on the disk.

#### **Server Disk Model**

The model name or model number of the resource.

#### **Server Disk Vendor**

The vendor who supplied the resource.

#### **Related tasks:**

“Running the Server Disks Capacity report” on page 668

To see information about the capacity of disks on servers, run the Server Disks Capacity report.

## **Predefined reports about file systems on servers**

You can run predefined reports to analyze the capacity of file systems that are grouped by server.

### **Running the File Systems Capacity report**

To see information about the capacity of file systems on servers, run the File Systems Capacity report.

To run the report with the default settings, click **Finish**.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **File Systems**.
4. Click **File Systems Capacity**.
5. Select one or more servers.
6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“File Systems Capacity report”

Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

### **File Systems Capacity report**

Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

### **Charts**

The bar chart shows the capacity, used space, available space, and percentage of used space for file systems on servers.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

For each server, the following information is provided:

#### **File System Mount Point**

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be /opt or /export/home.

#### **File System Type**

The type of file system that the resource uses.

#### **File System Used Space (GiB)**

The amount of used storage space in the file system of the resource.

#### **File System Available Space (GiB)**

The amount of unused storage space in the file system of the resource.

#### **File System Capacity (GiB)**

The amount of storage space on the file system of the resource.

#### **File System Used Space (%)**

The percentage of used space in the file system on the server.

**File System Available Inodes**

The number of unused inodes in file systems on the operating system.

**File System Used Inodes**

The number of used inodes in file systems on the operating system.

**Tip:** For Microsoft Windows systems, this property is blank.

**Related tasks:**

“Running the File Systems Capacity report” on page 669

To see information about the capacity of file systems on servers, run the File Systems Capacity report.

---

## Predefined reports about storage systems and components

You can run predefined reports to view capacity details and performance metrics for storage systems and their components. The components include controllers, modules, nodes, disks, I/O groups, managed disks, storage pools, ports, RAID arrays, and volumes.

### Predefined reports about storage systems

You can run predefined reports to view capacity details and performance metrics for storage systems.

**Running the Storage Systems Capacity report**

To see information about the capacity of disks, pools, and volumes on storage systems, run the Storage Systems Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Storage Systems Capacity**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Storage Systems Capacity report”

Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

**Storage Systems Capacity report**

Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

## Charts

The bar chart shows the used space and available space for pools on storage systems.

## Report output

For each storage system, the following information is provided:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

### Storage System Type

The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

### Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

### Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

### Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

### Storage System Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

### Storage System Physical Allocation Percentage

The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated pool space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

### Storage System Virtual Allocation Percentage

The percentage of physical space in storage system pools that is committed

to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

#### **Storage System Shortfall Percentage**

The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB  $\div$  100 GiB  $\times$  100).

This value is only available for pools with thin-provisioned volumes.

#### **Storage System Volume Capacity (GiB)**

The amount of space on all volumes on the storage system.

#### **Storage System Number of Volumes**

The number of volumes on a resource.

#### **Storage System Number of Disks**

The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

**Storage System Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage System Unallocated Disk Space (GiB)**

The amount of disk space that can be added to a pool.

**Storage System Volume Capacity Assigned to MDisks (GiB)**

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

**Storage System Disk Capacity (GiB)**

The total amount of space that is on a physical disk on the storage system.

**Storage System Total Disk Capacity (GiB)**

The amount of space on physical disks on a storage system, including spare disks.

**Storage System Cache (GiB)**

The size of the cache on the storage system. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage System Volume Capacity for z/OS**

The amount of space on all volumes on the storage system that the z/OS operating system can use.

**Storage System Real Available Pool Space (GiB)**

The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Real Configured Pool Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Related reports**

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Storage Systems Historical Capacity
- Disks Capacity
- Volumes Capacity
- Pools Capacity

**Related tasks:**

“Running the Storage Systems Capacity report” on page 671

To see information about the capacity of disks, pools, and volumes on storage systems, run the Storage Systems Capacity report.



## Running the Storage Systems Historical Capacity report

To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Storage Systems Historical Capacity**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Storage Systems Historical Capacity report” on page 563

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

## Storage Systems Historical Capacity report

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

## Charts

The chart shows the following space statistics for the storage systems:

- The capacity of the pools on the storage systems
- The used space in the pools on the storage systems
- The available space in the pools on the storage systems
- The capacity of the volumes on the storage systems
- The available space on the managed disks on the storage systems
- The capacity of the managed disks on the storage systems

## Report output

The following information is included in reports for storage systems:

### Storage System Name

A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Date** The date that the data was collected for the report.

**Storage System Pool Capacity (GiB)**

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

**Storage System Used Pool Space (GiB)**

The amount of space that is in use in all pools on a storage system.

**Storage System Available Pool Space (GiB)**

The amount of unused space in pools that are on the storage system.

**Storage System Volume Capacity (GiB)**

The amount of space on all volumes on the storage system.

**Storage System Number of Volumes**

The number of volumes on a resource.

**Storage System Number of Disks**

The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

**Storage System Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage System Unallocated Disk Space (GiB)**

The amount of disk space that can be added to a pool.

**Storage System Volume Capacity Assigned to MDisks (GiB)**

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

**Storage System Disk Capacity (GiB)**

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Total Disk Capacity (GiB)**

The amount of space on physical disks on a storage system, including spare disks.

**Storage System Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage System Volume Capacity for z/OS (GiB)**

The amount of space on all volumes on the storage system that the z/OS operating system can use.

**Storage System Real Available Pool Space (GiB)**

The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Real Configured Pool Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Related tasks:**

“Running the Storage Systems Historical Capacity report” on page 562  
To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

**Running the Most Active Storage Systems report**

To see which storage systems are the most active, run the Most Active Storage Systems report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Most Active Storage Systems**.
4. Optional: Select the configuration of the storage systems.
5. Select the storage systems, and then select a sort order.
6. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Storage Systems report” on page 678  
Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

## Most Active Storage Systems report

Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

### Charts

The bar chart shows the most active storage systems. By default, storage systems are sorted by the total I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

The information that is provided about the most active storage systems is as follows:

#### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

#### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

#### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

#### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Most Active Host Connections

- Most Active Nodes
- Most Active IO Groups
- Most Active Disks
- Most Active Managed Disks
- Most Active Pools
- Most Active Ports
- Most Active Volumes
- Most Active Controllers or Modules
- Most Active RAID Arrays
- Performance of One Storage System
- Storage Resource Relationships Summary

**Related tasks:**

“Running the Most Active Storage Systems report” on page 677

To see which storage systems are the most active, run the Most Active Storage Systems report.

### **Running the Performance of One Storage System report**

To see performance metrics for a storage system, run the Performance of One Storage System report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Performance of One Storage System**.
4. Optional: Select the configuration of the storage system.
5. Select a storage system.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Storage System report”

Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

### **Performance of One Storage System report**

Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

## Charts

The charts show the following performance metrics for a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, write I/O rates, and response times
- The read data rates, and the write data rates
- The read cache hit percentages, and the write cache hit percentages

## Report output

The following information is provided in the report for the period that you specified:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

### Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations.

### Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Related reports**

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Most Active Host Connections
- Most Active Nodes
- Most Active IO Groups
- Most Active Disks
- Most Active Managed Disks
- Most Active Pools
- Most Active Ports
- Most Active Volumes
- Most Active Controllers or Modules
- Most Active RAID Arrays

**Related tasks:**

“Running the Performance of One Storage System report” on page 679

To see performance metrics for a storage system, run the Performance of One Storage System report.

**Running the Compare Performance of Multiple Storage Systems report**

Use the Compare Performance of Multiple Storage Systems report to compare up to four performance metrics on multiple storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Compare Performance of Multiple Storage Systems**.
4. Select the storage systems, and then click **Next**.
5. Select the category of performance metric that you want to include in the report, and then select the performance metric.
6. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.

7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

9. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple Storage Systems report”

Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

### **Compare Performance of Multiple Storage Systems report**

Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

### **Charts**

The charts show the performance metrics for the storage systems that you selected for the period that you specified.

### **Report output**

The following information is provided in the report:

**Performance metrics**

Shows the data for each performance metric.

**Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

**Related tasks:**

“Running the Compare Performance of Multiple Storage Systems report” on page 681

Use the Compare Performance of Multiple Storage Systems report to compare up to four performance metrics on multiple storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

### **Running the Compare Performance of One Storage System over Time Ranges report**

To compare a performance metric for one storage system in two date ranges, run the Compare Performance of One Storage System over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Compare Performance of One Storage System over Time Ranges**.
4. Optional: Select the configuration of the storage systems.
5. Select the storage system.



6. Select the category of performance metric about which you want a report, and then select the performance metric.
7. Select an interval.
8. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

**Tip:** To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.

9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Compare Performance of One Storage System over Time Ranges report”

Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

## **Compare Performance of One Storage System over Time Ranges report**

Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

### **Charts**

The chart shows the performance data of the storage system on the performance metric you selected, over the time periods that you selected.

### **Report output**

The following information is provided in the report:

#### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### **Day in Range**

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

#### **Hour in Range**

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

#### **Minute in Range**

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5

minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

#### **First Date Range**

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

#### **Second Date Range**

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

#### **Related tasks:**

“Running the Compare Performance of One Storage System over Time Ranges report” on page 682

To compare a performance metric for one storage system in two date ranges, run the Compare Performance of One Storage System over Time Ranges report.

### **Running the Performance Data Export report**

To get a spreadsheet of the performance metrics for one or more resources on a storage system, run the Performance Data Export report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Performance Data Export**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Optional: Click **Next**, and then select the devices for which you want performance metrics in the report.
7. Click **Finish**.

You can open or save the spreadsheet from your browser.

#### **Related reference:**

“Performance Data Export report”

Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM.

### **Performance Data Export report**

Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM.

#### **Report output**

The spreadsheet contains a worksheet for each resource on the storage system. The worksheet for a resource contains the performance metric data for that resource.

#### **Related tasks:**

“Running the Performance Data Export report” on page 684

To get a spreadsheet of the performance metrics for one or more resources on a storage system, run the Performance Data Export report.

## Predefined reports about controllers, modules, and nodes

You can run predefined reports to analyze and compare the performance of controllers, modules, and nodes.

### Running the Most Active Controllers or Modules report

To see which controllers or modules are most active on storage systems, run the Most Active Controllers or Modules report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Most Active Controllers or Modules**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Most Active Controllers or Modules report”

Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

### Most Active Controllers or Modules report

Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

**Restriction:** This report applies only to DS8000, DS5000, XIV systems, and vendor-acquired systems.

### Charts

The bar chart shows the most active controllers or modules on storage systems. By default, controllers or modules are sorted by the overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

## Report output

The information that is provided about the most active controllers or modules is as follows:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Controller or Module Name

The name that was assigned to the controller or module when it was added to the system.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is *Storage System*, the storage system is configured as back-end storage. If this value is *Storage Virtualizer*, the storage system is configured to virtualize back-end storage.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Controller or module	Performance of One Controller or Module

### Related tasks:

“Running the Most Active Controllers or Modules report” on page 685

To see which controllers or modules are most active on storage systems, run the Most Active Controllers or Modules report.

### Running the Performance of One Controller or Module report

To see performance metrics for controllers or modules on a storage system, run the Performance of One Controller or Module report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Performance of One Controller or Module**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a controller or module.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Controller or Module report”

Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

### **Performance of One Controller or Module report**

Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

**Restriction:** This report applies only to DS8000, DS5000, XIV systems, and vendor-acquired systems.

### **Charts**

The charts show the following performance metrics for each controller or module on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, and the response times
- The read data rates, and the write data rates
- The overall cache hit percentage for read operations and write operations, and the utilization percentage for processors

### **Report output**

The following information is provided in the report for the period that you specified:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Controller Name or Module Name**

The name that was assigned to the controller or module when it was added to the system.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is *Storage System*, the storage system is configured as back-end storage. If this value is *Storage Virtualizer*, the storage system is configured to virtualize back-end storage.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Related reports**

To open the *Most Active Controllers or Modules* report, click the name of the controller or module in the report table.

**Related tasks:**

“Running the Performance of One Controller or Module report” on page 686  
To see performance metrics for controllers or modules on a storage system, run the Performance of One Controller or Module report.

## Running the Most Active Nodes report

To see which nodes are most active on your storage systems, run the Most Active Nodes report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Most Active Nodes**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Most Active Nodes report”

Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

## Most Active Nodes report

Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

## Charts

The chart shows the most active nodes on storage systems. By default, nodes are sorted by the total overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

## Report output

The information that is provided for nodes is as follows:

### Storage System Name

A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage I/O Group Name**

The name that was assigned to the I/O group when it was added to the system.

**Storage Node Name**

The name that was assigned to the storage system node when it was added to the system.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
I/O group	Performance of One IO Group
Node	Performance of One Node

**Related tasks:**

“Running the Most Active Nodes report” on page 689

To see which nodes are most active on your storage systems, run the Most Active Nodes report.

**Running the Performance of One Node report**

To see information about the performance of a node on a storage system, run the Performance of One Node report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Performance of One Node**.



5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a node.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Node report”

Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

### **Performance of One Node report**

Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### **Charts**

The charts show the following performance statistics for a node on a storage system:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

### **Report output**

The following information is provided in the report for the period that you specified:

#### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### **Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

**Storage I/O Group Name**

The name that was assigned to the I/O group when it was added to the system.

**Storage Node Name**

The name that was assigned to the storage system node when it was added to the system.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Most Active Nodes
I/O group	Performance of One IO Group

### Related tasks:

“Running the Performance of One Node report” on page 690

To see information about the performance of a node on a storage system, run the Performance of One Node report.

## Running the Compare Performance of Multiple Controllers and Modules report

Use this report to compare up to four performance metrics for multiple controllers and modules that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Compare Performance of Multiple Controllers and Modules**.
5. Select the storage systems, and then click **Refresh**.
6. Select the controllers and modules, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

### Related reference:

“Compare Performance of Multiple Controllers and Modules report”

Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

## **Compare Performance of Multiple Controllers and Modules report**

Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

**Restriction:** This report applies only to DS8000, DS5000, XIV systems, and vendor-acquired systems.

### **Charts**

The charts show the performance metrics for the controllers and modules that you selected for the period that you specified.

### **Report output**

The following information is provided in the report:

#### **Performance metric**

Shows the time at which the performance metric data was collected.

#### **Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

#### **Controller and modules**

Shows the names of the controllers and modules, and the performance metric data for each controller and module.

#### **Related tasks:**

“Running the Compare Performance of Multiple Controllers and Modules report” on page 693

Use this report to compare up to four performance metrics for multiple controllers and modules that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

## **Running the Compare Performance of Multiple Nodes report**

Use the Compare Performance of Multiple Controllers and Modules report to compare up to four performance metrics on multiple nodes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Compare Performance of Multiple Nodes**.
5. Select the storage systems, and then click **Refresh**.
6. Select the nodes, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:

- a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
  10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.
 

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
  11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple Nodes report”

Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

**Compare Performance of Multiple Nodes report**

Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

**Charts**

The charts show the performance metrics for the nodes that you selected for the period that you specified.

**Report output**

The following information is provided in the report:

**Performance metrics**

Shows the data for each performance metric.

**Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

**Nodes** Shows the names of the nodes, and the performance metric data for each node.

**Related tasks:**

“Running the Compare Performance of Multiple Nodes report” on page 694

Use the Compare Performance of Multiple Controllers and Modules report to compare up to four performance metrics on multiple nodes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

## Predefined reports about disks

You can run predefined reports to analyze the capacity of disks and performance of local disks.

### Running the Disks Capacity report

Run the Disks Capacity report to see information about the capacity of disks on storage systems.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Disks Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Disks Capacity report”

Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

### Disks Capacity report

Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

### Charts

The bar chart shows the capacity of disks, grouped by storage class, on storage systems.

### Report output

For each disk on a storage system, the following information is provided:

#### Storage Disk Class

The storage technology of the disk. This property applies only to DS8000, DS5000, and DS4000 storage systems. This property can contain the following values:

**FC** The disk is a Fibre Channel disk.

#### Nearline ATA

The disk is a nearline Advanced Technology Attachment disk.

#### Nearline SAS

The disk is a nearline serial-attached SCSI disk.

**SAS** The disk is a serial-attached SCSI disk.

**SATA** The disk is a Serial Advanced Technology Attachment disk.

**SSD** The disk is a solid-state drive.

**Unavailable**

The disk is on a DS5000 or DS4000 series system, and is not a solid-state drive.

**Storage Disk Name**

The name that was assigned to the storage disk when it was added to the system.

**Storage System Type**

The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Disk Tag**

A number that identifies the array to which a disk belongs. This property applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems.

**Storage Disk Vendor**

The vendor who supplied the resource.

**Storage Disk Model**

The model name or model number of the resource.

**Storage Disk Serial Number**

The serial number of the resource.

**Storage Disk Firmware Version**

The version number of the firmware that is running on the disk.

**Storage Disk Capacity (GiB)**

The amount of storage space that is on a physical disk on the storage system.

**Storage Disk Speed (RPM)**

The speed of the disk.

**Related tasks:**

“Running the Disks Capacity report” on page 696

Run the Disks Capacity report to see information about the capacity of disks on storage systems.

## Running the Most Active Disks report

To see which local disks that are on storage systems are the most active, use the Most Active Disks report. This report shows information only for SAN Volume Controller and Storwize V7000 storage systems.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Most Active Disks**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Most Active Disks report”

Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

## Most Active Disks report

Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

## Charts

The bar chart shows the most active local disks that are on storage systems. By default, local disks are sorted by the total I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.



## Report output

The information that is provided about the most active local disks that are on storage systems is as follows:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage MDisk Name

The name that was assigned to the managed disk when it was added to the system.

### Storage Local Disk Name

The name that was assigned to the local disk on the storage system when the local disk was created on the system.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### Storage Local Disk Storage Class

The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

### Storage Local Disk Speed (RPM)

The speed of the local disk on a storage system.

### Storage Local Disk Capacity (GiB)

The amount of storage space on the local disk.

### Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System

Resource	Related reports
Managed disk	Performance of One Managed Disk
Disk	Performance of One Disk

**Related tasks:**

“Running the Most Active Disks report” on page 698

To see which local disks that are on storage systems are the most active, use the Most Active Disks report. This report shows information only for SAN Volume Controller and Storwize V7000 storage systems.

**Running the Performance of One Disk report**

To see performance metrics for a local disk on a storage system, use the Performance of One Disk report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Performance of One Disk**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a local disk.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Disk report”

Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

**Performance of One Disk report**

Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

**Charts**

The charts show the following performance metrics for a local disk on a storage system:

- The total back-end I/O rates and the overall back-end response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times

- The read data rates, the write data rates, and the total data rates
- The read queue times, the write queue times, and the overall queue times

## Report output

The following information is provided in the report for the period that you specified:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### Storage MDisk Name

The name that was assigned to the managed disk when it was added to the system.

### Storage Local Disk Name

The name that was assigned to the local disk on the storage system when the local disk was created on the system.

### Storage Local Disk Storage Class

The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

### Storage Local Disk Speed (RPM)

The speed of the local disk on a storage system.

### Storage Local Disk Capacity (GiB)

The amount of storage space on the local disk.

### Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

### Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Related reports**

To open the Most Active Disks report, click the name of the storage system in the report table.

**Related tasks:**

“Running the Performance of One Disk report” on page 700

To see performance metrics for a local disk on a storage system, use the Performance of One Disk report.

**Running the Compare Performance of Multiple Disks report**

Use the Compare Performance of Multiple Disks report to compare up to four performance metrics for multiple local disks that are on storage systems. You can view reports about local disks on SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and Storwize V3700 systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Compare Performance of Multiple Disks**.
5. Select the storage systems, and then click **Refresh**.
6. Select the local disks, and then click **Next**.

7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple Disks report”

Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

**Compare Performance of Multiple Disks report**

Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

**Charts**

The charts show the performance metrics for the local disks that you selected for the period that you specified.

**Report output**

The following information is provided in the report:

**Performance metric**

Shows the time at which the performance metric data was collected.

**Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

**Disks** Shows the names of the disks, and the performance metric data for each disk.

**Related tasks:**

“Running the Compare Performance of Multiple Disks report” on page 702  
Use the Compare Performance of Multiple Disks report to compare up to four performance metrics for multiple local disks that are on storage systems. You can view reports about local disks on SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and Storwize V3700 systems. For example, you can compare cache percentages, data rates, and I/O rates.

## Predefined reports about host connections

You can run predefined reports to analyze and compare the performance of volumes on host connections.

### Running the Most Active Host Connections report

To see which host connections are the most active, run the Most Active Host Connections report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Host Connections**.
4. Click **Most Active Host Connections**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Most Active Host Connections report”

Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

### Most Active Host Connections report

Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

#### Chart

The bar chart shows the most active host connections. By default, host connections are sorted by the total I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

## Report output

The information that is provided about the most active host connections is as follows:

### Storage Host Connection Name

The user-defined name that describes the server or storage device that is assigned to a volume.

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### XIV Cluster

The name of the cluster, if any, as defined on the XIV system. If no cluster is defined for the host connection, this field is blank.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Host connection	Summarized Performance of Volumes by Host Connection
Storage system	Performance of One Storage System

### Related tasks:

“Running the Most Active Host Connections report” on page 704

To see which host connections are the most active, run the Most Active Host Connections report.

## Running the Summarized Performance of Volumes by Host Connection report

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Host Connection report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Host Connections**.
4. Click **Summarized Performance of Volumes by Host Connection**.
5. Select a host connection. To select a host connection, complete the following steps:
  - a. Type a keyword. For example, type `example1` to find the names of host connections that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find host connections that begin with 'a' or 'A'. Type `%a` to find host connections that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the host connections on a resource.
  - b. Click **Search**. The report searches the user-defined names for the host connections. The report does not search the server names that are defined when a server is probed. If the search is successful, the host connections are displayed in the **Choice** list.
  - c. From the list, select a host connection, and then click **Refresh**.
6. Select the volumes.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

"Summarized Performance of Volumes by Host Connection report"

Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

## Summarized Performance of Volumes by Host Connection report

Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

### Charts

The charts show the following performance statistics for the volumes on a host connection:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times



- The read data rates, the write data rates, and the total data rates
- The read cache percentages and the write cache percentages

## Report output

The following information is provided in the report for the period that you specified:

### Storage Host Connection Name

The user-defined name that describes the server or storage device that is assigned to a volume.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation, across all volumes on the server. This value includes both read and write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second, across all volumes on the server. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the server.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the server. The average value for each volume is added to give a total data rate for all volumes on the server.

### Write Response Time (ms/op)

The average number of milliseconds to complete a write operation, across all volumes on the server.

### Read Response Time (ms/op)

The average number of milliseconds to complete a read operation, across all volumes on the server.

### Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations, across all volumes on the server. The average value for each volume is added to give a read data rate for all volumes on the server.

### Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations, across all volumes on the server. The average value for each volume is added to give a write data rate for all volumes on the server.

### Overall Read I/O Rate (ops/s)

The average number of read operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential read operations.

### Overall Write I/O Rate (ops/s)

The average number of write operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential write operations.

### Overall Read Cache Hit Percentage

The average percentage of all read operations that find data in the cache, across all volumes on the server. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### Overall Write Cache Hit Percentage

The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

### Related reports

To open the Performance of Volumes by Host Connection report, click the name of the host connection in the report table.

#### Related tasks:

“Running the Summarized Performance of Volumes by Host Connection report” on page 706

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Host Connection report.

### Running the Performance of Volumes by Host Connection report

To see performance metrics for the volumes on a server, use the Performance of Volumes by Host Connection report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Host Connections**.
4. Click **Performance of Volumes by Host Connection**.
5. Select a host connection. To select a host connection, complete the following steps:
  - a. Type a keyword. For example, type `example1` to find the names of host connections that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find host connections that begin with 'a' or 'A'. Type `%a` to find host connections that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the host connections on a resource.
  - b. Click **Search**. The report searches the user-defined names for the host connections. The report does not search the server names that are defined when a server is probed. If the search is successful, the host connections are displayed in the **Choice** list.
  - c. From the list, select a host connection, and then click **Refresh**.
6. Select the volumes.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Performance of Volumes by Host Connection report”

Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

### **Performance of Volumes by Host Connection report**

Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

### **Charts**

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. This is a stacked area chart that shows the contribution of each volume to the overall load on the server.
- The total data rates. This is a stacked area chart that shows the contribution of each volume to the overall load on the server.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

### **Report output**

The following information is provided in the report for the period that you specified:

#### **Storage Host Connection Name**

The user-defined name that describes the server or storage device that is assigned to a volume.

#### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### **Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

#### **Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

#### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

#### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### **Volume Utilization**

The average percentage of time that the volume is busy.

#### **Related tasks:**

“Running the Performance of Volumes by Host Connection report” on page 708  
To see performance metrics for the volumes on a server, use the Performance of Volumes by Host Connection report.

## Predefined reports about I/O groups

You can run predefined reports to analyze and compare the performance of I/O groups.

### Running the Most Active IO Groups report

To see which I/O groups are most active on storage systems, run the Most Active IO Groups report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **IO Groups**.
4. Click **Most Active IO Groups**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Most Active IO Groups report”

Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

### Most Active IO Groups report

Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

## Charts

The chart shows the most active I/O groups on storage systems. By default, I/O groups are sorted by the total overall I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

## Report output

The information that is provided for I/O groups is as follows:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

### Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
I/O group	Performance of One IO Group

### Related tasks:

“Running the Most Active IO Groups report” on page 710

To see which I/O groups are most active on storage systems, run the Most Active IO Groups report.

## Running the Performance of One IO Group report

To see information about the performance of an I/O group on a storage system, run the Performance of One IO Group report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **IO Groups**.
4. Click **Performance of One IO Group**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select an I/O group.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Performance of One IO Group report”

Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

## Performance of One IO Group report

Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Charts

The charts show the following performance metrics for an I/O group on a storage system:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

### Report output

The following information is provided in the report for the period that you specified:

#### Storage System Name

A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Storage I/O Group Name**

The name that was assigned to the I/O group when it was added to the system.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

## Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy.

## Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Performance of One Storage System
- Most Active IO Groups

### Related tasks:

“Running the Performance of One IO Group report” on page 712

To see information about the performance of an I/O group on a storage system, run the Performance of One IO Group report.

## Running the Compare Performance of Multiple IO Groups report

Use the Compare Performance of Multiple IO Groups report to compare up to four performance metrics for multiple I/O groups that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **IO Groups**.
4. Click **Compare Performance of Multiple IO Groups**.
5. Select the storage systems, and then click **Refresh**.
6. Select the I/O groups, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

### Related reference:

“Compare Performance of Multiple IO Groups report” on page 715

Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.



## Compare Performance of Multiple IO Groups report

Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Charts

The charts show the performance metrics for the I/O groups that you selected for the period that you specified.

### Report output

The following information is provided in the report:

#### Performance metric

Shows the time at which the performance metric data was collected.

#### Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

#### I/O groups

Shows the names of the I/O groups, and the performance metric data for each I/O group.

#### Related tasks:

“Running the Compare Performance of Multiple IO Groups report” on page 714  
Use the Compare Performance of Multiple IO Groups report to compare up to four performance metrics for multiple I/O groups that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

## Predefined reports about managed disks

You can run predefined reports about to analyze and compare the capacity and performance of managed disks.

### Running the Managed Disks Capacity report

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Managed Disks Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Managed Disks Capacity report” on page 521

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

### **Managed Disks Capacity report**

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

**Restriction:** This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

### **Charts**

The bar chart shows the used space and the available space for managed disk on storage systems.

### **Report output**

For each managed disk on a storage system, the following information is provided:

**Storage System MDisk Name**

The name that was assigned to the managed disk on a storage system when it was added to the system.

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage System MDisk Used Space (GiB)**

The amount of used space on the managed disk.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage MDisk RAID Level**

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

**Storage System MDisk Type**

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

**Storage System MDisk Mode**

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

**Array** Extents are provided from local disks.

**Managed**

Extents are provided from a back-end storage volume.

**Unmanaged**

The managed disk is not used in the system.

**Storage System MDisk Strip Size (KB)**

The RAID strip size on a managed disk on a storage system.

**Storage System MDisk Is Balanced**

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is Yes, the LUNs are balanced.

**Storage System MDisk Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

**Storage System MDisk Write Verify**

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is Yes, all write operations are verified by a follow-up read operation.

**Related tasks:**

“Running the Managed Disks Capacity report” on page 521

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

**Running the Most Active Managed Disks report**

To see which managed disks are most active on storage systems, use the Most Active Managed Disks.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Most Active Managed Disks**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Managed Disks report”

Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

### **Most Active Managed Disks report**

Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

**Restriction:** This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

### **Charts**

The bar chart shows the most active managed disks that are on storage systems. By default, managed disks are sorted by the I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided about the most active managed disks is as follows:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage MDisk Name**

The name that was assigned to the managed disk when it was added to the system.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool
Managed disk	Performance of One Managed Disk, Storage Resource Relationships Summary

**Related tasks:**

“Running the Most Active Managed Disks report” on page 717

To see which managed disks are most active on storage systems, use the Most Active Managed Disks.

**Running the Performance of One Managed Disk**

To see performance metrics for managed disks on a storage system, use the Performance of One Managed Disk.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Performance of One Managed Disk**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a managed disk.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Managed Disk report” on page 720

Shows four charts, and a table of performance metrics for a managed disk over a

time period. Use the report to analyze the performance of a managed disk.

### **Performance of One Managed Disk report**

Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

**Restriction:** This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

### **Charts**

The charts show the following performance statistics for a managed disk on a storage system:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates and the write data rates
- The read cache percentages, the write cache percentages, and the volume utilization percentages

### **Report output**

The following information is provided in the report for the period that you specified:

#### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### **Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is *Storage System*, the storage system is configured as back-end storage. If this value is *Storage Virtualizer*, the storage system is configured to virtualize back-end storage.

#### **Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

#### **Storage MDisk Name**

The name that was assigned to the managed disk when it was added to the system.

#### **Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

#### **Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

#### **Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Most Active Managed Disks
Storage pool	Performance of One Pool

**Related tasks:**

“Running the Performance of One Managed Disk” on page 719

To see performance metrics for managed disks on a storage system, use the Performance of One Managed Disk.

## Running the Compare Performance of Multiple Managed Disks report

Use the Compare Performance of Multiple Managed Disks report to compare up to four performance metrics for multiple managed disks that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Compare Performance of Multiple Managed Disks**.
5. Select the storage systems.
6. Select managed disks. To select managed disks, complete the following steps:
  - a. Type a keyword. For example, type `mdisk1` to find the names of managed disks that start with `mdisk1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find managed disks that begin with 'a' or 'A'. Type `%a` to find managed disks that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the managed disks on a resource.
  - b. Click **Search**. If the search is successful, the managed disks are displayed in the **Results** list.
  - c. Select the managed disks that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
12. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

### Related reference:

"Compare Performance of Multiple Managed Disks report" on page 723  
Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.



## Compare Performance of Multiple Managed Disks report

Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

**Restriction:** This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

### Charts

The charts show the performance metrics for the managed disks that you selected for the period that you specified.

### Report output

The following information is provided in the report:

#### Performance metric

Shows the time at which the performance metric data was collected.

#### Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

#### Managed disks

Shows the names of the managed disks, and the performance metric data for each managed disk.

#### Related tasks:

“Running the Compare Performance of Multiple Managed Disks report” on page 722

Use the Compare Performance of Multiple Managed Disks report to compare up to four performance metrics for multiple managed disks that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

## Predefined reports about storage pools

You can run predefined reports to analyze and compare the capacity and performance of storage pools.

### Running the Pools Capacity report

To see information about the used and available space of pools on storage systems, run the Pools Capacity report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Pools Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Optional: Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Pools Capacity report”

Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

**Pools Capacity report**

Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

**Charts**

The bar chart shows the used space and available space for pools on storage systems.

**Report output**

For each pool on a storage system, the following information is provided:

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Storage Pool Capacity (GiB)**

The total amount of storage space in a pool.

**Storage Pool Used Space (GiB)**

The amount of storage space that is used in the pool.

**Storage Pool Available Space (GiB)**

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{pool capacity} - \text{used space}$$

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

**Storage Pool Physical Allocation Percentage**

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

### **Storage Pool Virtual Allocation Percentage**

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

### **Storage Pool Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{virtual capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the

committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

**Storage Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Pool Real Configured Space (GiB)**

The amount of storage space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage Pool Real Available Space (GiB)**

The amount of unused space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage Pool RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Pool Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Pool Is Solid State**

Shows whether there are solid-state drives in the pool. This property can contain the following values:

**Mixed** The pool contains both hard disk drives and solid-state drives.

**Non solid state**

The pool contains no solid-state drives.

**Solid state**

The pool contains a least one solid-state drive.

**Storage Pool Number of MDisks**

The number of managed disks in a storage pool.

**Storage Pool Extent Size (MiB)**

The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage Pool Warning Level**

The percentage of used capacity of the storage pool at which a warning is generated.

**Storage Pool Repository Capacity (GiB)**

The amount of space on all extents in the repository of a pool. This space can be used to allocate track space-efficient volumes. This attribute applies only to DS8000 storage systems.

**Storage Pool Used Repository Space (GiB)**

The amount of space on all extents in the repository of a pool that are allocated. This space can be used to allocate track space-efficient volumes. This attribute applies only to the DS8000 storage systems.

**Report links**

To open related reports, click the name of the storage pool and then click the name of the report. You can open the following related reports:

- Pools Historical Capacity
- Volumes Capacity

**Related tasks:**

“Running the Pools Capacity report” on page 723

To see information about the used and available space of pools on storage systems, run the Pools Capacity report.

**Running the Pools Historical Capacity report**

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

You can see how space in a pool on a storage system, or in pools in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Pools Historical Capacity**.
5. Select the type of historical space report that you want to create. You can select either of the following options:
  - **Of a Pool in a Storage System**
    - a. Click **Next**.
    - b. Optional: Select the configuration of the storage system.
    - c. Select a storage system, and then select a pool.
  - **Of Pools in a Storage Resource Group**
    - a. Click **Next**.
    - b. Optional: Select the configuration of the storage system.
    - c. Select a storage resource group.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Pools Historical Capacity report” on page 565

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

### **Pools Historical Capacity report**

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

### **Charts**

The chart shows the following space statistics for the pool on a storage system:

- The capacity of the pool
- The used space in the pool
- The available space in the pool
- The virtual volume capacity of the pool
- The used volume space in the pool
- The allocated space in the pool

The chart shows the following space statistics for the pool in the storage resource group:

- The capacity of the pools

- The used space in the pools
- The available space in the pools
- The virtual volume capacity of the pools
- The used volume space in the pools
- The allocated space in the pool

## Report output

The following information is included in reports for a pool on a storage system or for pools in a storage resource group:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### Storage Pool Name

The name that was assigned to the pool when it was added to the system.

**Date** The date that the data was collected for the report.

### Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

### Storage Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{pool capacity} - \text{used space}$$

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

### Storage Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage

- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Pool Used Space (GiB)**

The amount of storage space that is used in the pool.

**Storage Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Pool Real Configured Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage Pool Number of MDisks**

The number of managed disks in a storage pool.

**Storage Pool Number of Storage Volumes**

The number of volumes in the storage pool.

**Related tasks:**

“Running the Pools Historical Capacity report” on page 564

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

**Running the Most Active Pools report**

To see which pools are most active on storage systems, run the Most Active Pools report.



**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Most Active Pools**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Pools report”

Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

### **Most Active Pools report**

Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

### **Charts**

The bar chart shows the most active pools on storage systems. By default, pools are sorted by the I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided about the most active pools is as follows:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System,

the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool, Most Active Volumes, Most Active RAID Arrays, Storage Resource Relationships Summary

**Related tasks:**

“Running the Most Active Pools report” on page 730

To see which pools are most active on storage systems, run the Most Active Pools report.

**Running the Performance of One Pool report**

To see information about the performance of a pool on a storage system, run the Performance of One Pool report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Performance of One Pool**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a pool.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Pool report”

Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

### **Performance of One Pool report**

Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

### **Charts**

The charts show the following performance metrics for a pool on a storage system:

- The total overall rates for I/O operations, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The total I/O rates for back-end operations, and overall response times for back-end operations
- The I/O rates for back-end read operations and back-end write operations, and the response rates for back-end read and back-end write operations

### **Report output**

The following information is provided in the report for the period that you specified:

#### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### **Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

#### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

#### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

#### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

#### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

#### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### **Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Related reports**

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Performance of One Storage System
- Most Active Pools

**Related tasks:**

“Running the Performance of One Pool report” on page 732

To see information about the performance of a pool on a storage system, run the Performance of One Pool report.

**Running the Compare Performance of Multiple Pools**

Use the Compare Performance of Multiple Pools report to compare up to four performance metrics for multiple pools that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Compare Performance of Multiple Pools**.
5. Select the storage systems, and then click **Refresh**.
6. Select the pools, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.

- b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
  10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple Pools report”

Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

## **Compare Performance of Multiple Pools report**

Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

## **Charts**

The charts show the performance metrics for the storage pools that you selected for the period that you specified.

## **Report output**

The following information is provided in the report:

### **Day, hour, or time**

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

### **Performance metrics**

Shows the performance metric data for the resource.

### **Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

**Pools** Shows the names of the storage pools, and the performance metric data for each storage pool.

### **Related tasks:**

“Running the Compare Performance of Multiple Pools” on page 734

Use the Compare Performance of Multiple Pools report to compare up to four performance metrics for multiple pools that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

## Running the Compare Performance of One Pool over Time Ranges report

To compare a performance metric for one pool in two date ranges, run the Compare Performance of One Pool over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Compare Performance of One Pool over Time Ranges**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage system.
7. Select the category of performance metric about which you want a report, and then select the performance metric.
8. Select an interval.
9. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

**Tip:** To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Compare Performance of One Pool over Time Ranges report”

Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

## Compare Performance of One Pool over Time Ranges report

Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

### Charts

The chart shows the performance data of the pool on the performance metric you selected, over the time periods that you selected.

### Report output

The following information is provided in the report:

#### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### Storage Pool Name

The name that was assigned to the pool when it was added to the system.

#### Day in Range

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The

number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

#### **Hour in Range**

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

#### **Minute in Range**

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

#### **First Date Range**

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

#### **Second Date Range**

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

#### **Related tasks:**

“Running the Compare Performance of One Pool over Time Ranges report” on page 736

To compare a performance metric for one pool in two date ranges, run the Compare Performance of One Pool over Time Ranges report.

## **Predefined reports about ports**

You can run predefined reports to analyze and compare the performance of ports.

### **Running the Most Active Ports report**

To see which ports are most active on storage systems, run the Most Active Ports report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Ports**.
4. Click **Most Active Ports**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.

7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Ports report”

Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

### **Most Active Ports report**

Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

### **Charts**

The bar chart shows the most active ports on storage systems. By default, ports are sorted by the I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided about the most active ports is as follows:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Port Name**

The name that was assigned to the storage port when the storage system was added to the system.

**Storage System Port Location**

The physical location of a port on DS8000, DS6000, TotalStorage Enterprise Storage Server storage systems, and on XIV systems. For all other systems, this property shows the WWPN of the port.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Overall Port Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.



### Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

### Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

## Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Number of the Fibre Channel (FC) port	Performance of One Port

### Related tasks:

“Running the Most Active Ports report” on page 737

To see which ports are most active on storage systems, run the Most Active Ports report.

## Running the Performance of One Port report

To see performance metrics for a port on a storage system, run the Performance of One Port report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Ports**.
4. Click **Performance of One Port**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a port.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Performance of One Port report” on page 740

Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

## Performance of One Port report

Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

### Charts

The charts show the following performance metrics for a port on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, and the response times
- The read data rates, the write data rates, and the total data rate
- The send bandwidth percentage, the receive bandwidth percentage, and the overall bandwidth percentage

### Report output

The following information is provided in the report for the period that you specified:

#### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

#### Storage Port Name

The name that was assigned to the storage port when the storage system was added to the system.

#### Storage System Port Location

The physical location of a port on DS8000, DS6000, TotalStorage Enterprise Storage Server storage systems, and on XIV systems. For all other systems, this property shows the WWPN of the port.

#### Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port.

#### Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data.

#### Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

#### Port Send Response Time (ms/op)

The average number of milliseconds to complete a send operation.

#### Port Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation.

#### Overall Port Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Related reports**

To open the Most Active Ports report, click the name of the storage system in the report table.

**Related tasks:**

“Running the Performance of One Port report” on page 739

To see performance metrics for a port on a storage system, run the Performance of One Port report.

**Running the Compare Performance of Multiple Ports report**

Use the Compare Performance of Multiple Ports report to compare up to four performance metrics for multiple ports that are on storage systems. For example, you can compare data rates, I/O rates, and bandwidth percentages.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Ports**.
4. Click **Compare Performance of Multiple Ports**.
5. Select the storage systems.
6. Select ports. To select ports, complete the following steps:
  - a. Type a keyword. For example, type port1 to find the names of ports that start with port1. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type a% to find ports that begin with 'a' or 'A'. Type %a to find ports that begin with or contain the letter 'a' or 'A'. You can type % to retrieve all of the ports on a resource.
  - b. Click **Search**. If the search is successful, the ports are displayed in the **Results** list.
  - c. Select the ports that you require from the list, and then click **Insert**.

7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple Ports report”

Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

### **Compare Performance of Multiple Ports report**

Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

### **Charts**

The charts show the performance metrics for the ports that you selected for the period that you specified.

### **Report output**

The following information is provided in the report:

**Day, hour, or time**

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

**Performance metrics**

Shows the performance metric data for the resource.

**Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

**Ports** Shows the names of the storage ports, and the performance metric data for each port.

**Related tasks:**

“Running the Compare Performance of Multiple Ports report” on page 741  
Use the Compare Performance of Multiple Ports report to compare up to four performance metrics for multiple ports that are on storage systems. For example, you can compare data rates, I/O rates, and bandwidth percentages.

## Predefined reports about RAID arrays

You can run predefined reports to analyze and compare the performance of RAID arrays.

### Running the Most Active RAID Arrays report

To see which arrays are most active on DS8000, DS6000, and IBM TotalStorage Enterprise Storage Server storage systems, run the Most Active RAID Arrays report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **RAID Arrays**.
4. Click **Most Active RAID Arrays**.
5. Select the storage systems, and then select a sort order.
6. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

#### Related reference:

“Most Active RAID Arrays report”

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

### Most Active RAID Arrays report

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

**Restriction:** This report applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server systems.

### Charts

The bar chart shows the most active RAID on DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems. By default, arrays are sorted by the I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### Report output

The information that is provided about the most active RAID is as follows:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Controller Name**

The name that was assigned to the controller when it was added to the system.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage RAID Array Name**

The name that was assigned to the RAID array when it was added to the system.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Disk Utilization Percentage**

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed for this property. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

**Related reports**

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Controller	Performance of One Controller or Module
RAID array	Performance of One RAID Array

**Related tasks:**

“Running the Most Active RAID Arrays report” on page 743

To see which arrays are most active on DS8000, DS6000, and IBM TotalStorage Enterprise Storage Server storage systems, run the Most Active RAID Arrays report.

**Running the Performance of One RAID Array report**

To see performance metrics for a RAID array on DS8000, DS6000, or IBM TotalStorage Enterprise Storage Server storage systems, run the Performance of One RAID Array report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **RAID Arrays**.
4. Click **Performance of One RAID Array**.
5. Select a storage system.
6. Select a RAID array.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One RAID Array report”

Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

**Performance of One RAID Array report**

Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

**Restriction:** This report applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server systems.

**Charts**

The charts show the following performance statistics for arrays on storage systems:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times

- The back-end I/O rates, and the back-end I/O response times
- The back-end read I/O rates, the back-end write I/O rates, the read response times, and the write response times

## Report output

The following information is provided in the report for the period that you specified:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage System Controller Name

The name that was assigned to the controller when it was added to the system.

### Storage RAID Array Name

The name that was assigned to the RAID array when it was added to the system.

### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

### Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

### Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

### Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

### Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.



**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Disk Utilization Percentage**

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed for this property. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

**Related reports**

To open the Most Active RAID Arrays report, click the name of the storage system in the report table.

**Related tasks:**

“Running the Performance of One RAID Array report” on page 745

To see performance metrics for a RAID array on DS8000, DS6000, or IBM TotalStorage Enterprise Storage Server storage systems, run the Performance of One RAID Array report.

**Running the Compare Performance of Multiple RAID Arrays report**

To compare the performance of multiple arrays that are on storage systems, use the Compare Performance of Multiple RAID Arrays report. For example, you can compare data rates, I/O rates, and response times.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **RAID Arrays**.
4. Click **Compare Performance of Multiple RAID Arrays**.
5. Select the storage systems, and then click **Refresh**.
6. Select the RAID, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.

9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple RAID Arrays report”

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

### **Compare Performance of Multiple RAID Arrays report**

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

**Restriction:** This report applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server systems.

### **Charts**

The charts show the performance metrics for the arrays that you selected for the period that you specified.

### **Report output**

The following information is provided in the report:

**Performance metric**

Shows the time at which the performance metric data was collected.

**Storage system names**

Shows the names of the storage systems, and the performance metric data for each storage system.

**RAID arrays**

Shows the names of the RAID arrays, and the performance metric data for each RAID array.

**Related tasks:**

“Running the Compare Performance of Multiple RAID Arrays report” on page 747  
To compare the performance of multiple arrays that are on storage systems, use the Compare Performance of Multiple RAID Arrays report. For example, you can compare data rates, I/O rates, and response times.

## **Predefined reports about volumes**

You can run predefined reports to analyze the capacity and performance of volumes.

## Running the Volumes Capacity report

To see information about the capacity of volumes on storage systems, run the Volumes Capacity report.

To run the report with the default settings, click **Finish**. By default all of the storage systems in the list are selected.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Volumes Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Optional: If you selected SAN Volume Controller or Storwize V7000 storage systems, you can select the type of volume.

Option	Description
All Volumes	Select <b>All</b> , the default, to view information about all of the volumes on the selected storage systems.
Only Primary Volumes	Select <b>Primary</b> to view information about primary volumes on the selected storage systems.
Only Secondary Volumes	Select <b>Secondary</b> to view information about secondary volumes on the selected storage systems.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

“Volumes Capacity report”

Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

## Volumes Capacity report

Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

## Charts

The bar chart shows the following statistics for volumes on storage systems:

- Used space.
- Unallocated space.
- Unused allocated space, that is, the value of Storage Volume Allocated Space (GiB) minus the value of Storage Volume Used Space (GiB).

## Report output

For each volume on a storage system, the following information is provided:

### **Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

### **Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### **Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

### **Storage Volume WWN**

The worldwide name of the volume.

### **Storage Volume Capacity (GiB)**

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

### **Storage Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

### **Storage Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This

inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

#### **Storage Volume Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine this value:

capacity - allocated space

Available only for thin provisioned volumes.

#### **Storage Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

#### **Storage Volume Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$(\text{allocated space} \div \text{capacity}) \times 100$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

#### **Storage Volume Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$(\text{unallocatable volume space} \div \text{volume unallocated space}) \times 100$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space} \div 100 \text{ GiB unallocated volume space} \times 100$

**Storage Volume Used Allocated Space Percentage**

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 volumes.

**Storage Volume RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Volume Is Assigned**

Shows whether the volume is assigned to a server. If this value is Yes, the volume is assigned to a server.

**Storage Volume Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Volume Type**

The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Warning Level**

The percentage of volume capacity that is used at which a warning is generated. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Logical Subsystem**

The logical subsystem (LSS) to which a volume or pool belongs.

**Storage Volume Copy ID**

The identifier of a volume copy.

**Storage Volume Copy Is Primary**

Shows whether the volume copy is the primary volume copy.

**Related reports**

To open the Volumes Historical Capacity report, click the name of the volume in the report table.

**Related tasks:**

“Running the Volumes Capacity report” on page 749

To see information about the capacity of volumes on storage systems, run the

Volumes Capacity report.

## Running the Volumes Historical Capacity report

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

You can see how space on a volume on a storage system, or on volumes in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Volumes Historical Capacity**.
5. Select the type of historical capacity report that you want to create. You can select either of the following options:
  - **Of a Volume in a Storage System**
    - a. Click **Next**.
    - b. Optional: Select the configuration of the storage system.
    - c. Select a storage system.
    - d. Type a keyword and click **Search**. For example, type `1un`, to find the names of volumes that start with `1un`. To change the default search option, click **Options**.

**Tip:** Use the percent sign (%) as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.

- e. Select a volume.
  - **Of Volumes in a Storage Resource Group**
    - a. Click **Next**.
    - b. Select a storage resource group.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

### Related reference:

"Volumes Historical Capacity report" on page 568

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

## Volumes Historical Capacity report

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

### Charts

The chart shows the following space statistics for the volume on a storage system:

- The capacity of the volume
- The used space on the volume
- The space on the volume that is used for system management
- The allocated space on the volume

The chart shows the following space statistics for the volumes in a storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The space on the volumes that is used for system management
- The allocated space on the volume

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

### Report output

The following information is included in reports for a volume on a storage system, or for volumes in a storage resource group:

#### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

#### Storage Pool Name

The name that was assigned to the pool when it was added to the system.

#### Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

**Date** The date that the data was collected for the report.

#### Storage Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.



For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

#### **Storage Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

#### **Storage Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

#### **Storage Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

#### **Related tasks:**

“Running the Volumes Historical Capacity report” on page 567

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

### **Running the Most Active Volumes report**

To see which volumes are most active on storage systems, run the Most Active Volumes report.

**Tip:** To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Most Active Volumes**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Most Active Volumes report”

Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

### **Most Active Volumes report**

Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

### **Charts**

The bar chart shows the most active volumes on storage systems. By default, volumes are sorted by the I/O rate.

**Tip:** To change the sort order, select a value from the **Sort Order** list.

### **Report output**

The information that is provided about the most active volumes is as follows:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### Volume Utilization

The average percentage of time that the volume is busy.

### Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool
Volume	Performance of One Volume, Storage Resource Relationships Summary

### Related tasks:

“Running the Most Active Volumes report” on page 755

To see which volumes are most active on storage systems, run the Most Active Volumes report.

### Running the Performance of One Volume report

To see performance metrics for volumes on a storage system, run the Performance of One Volume report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Performance of One Volume**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a volume. To select a volume, complete the following steps:
  - a. Type a keyword. For example, type lun, to find the names of volumes that start with lun. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type a% to find volume names that begin with 'a' or 'A'. Type %a to find volume names that begin with or contain the letter 'a' or 'A'. You can type % to retrieve all of the volumes on a resource.
  - b. Click **Search**. If the search is successful, the names of the volumes are displayed in the **Choice** list.
  - c. From the list, select a volume.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

**Related reference:**

“Performance of One Volume report”

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

### **Performance of One Volume report**

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

### **Charts**

The charts show the following performance metrics for a volume on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, and the write data rates
- The read cache percentages, the write cache percentages, and the volume utilization percentages

### **Report output**

The following information is provided in the report for the period that you specified:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Related reports**

To open the Most Active Volumes report, click the name of the storage system in the report table.

**Related tasks:**

“Running the Performance of One Volume report” on page 757

To see performance metrics for volumes on a storage system, run the Performance of One Volume report.

**Running the Compare Performance of Multiple Volumes report**

Use the Compare Performance of Multiple Volumes report to compare up to four performance metrics for multiple volumes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.

3. Click **Volumes**.
4. Click **Compare Performance of Multiple Volumes**.
5. Select the storage systems.
6. Select volumes. To select volumes, complete the following steps:
  - a. Type a keyword. For example, type `example1` to find the names of volumes that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find volumes that begin with 'a' or 'A'. Type `%a` to find volumes that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
  - b. Click **Search**. If the search is successful, the volumes are displayed in the **Results** list.
  - c. Select the volumes that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
  - a. Click **Add Metric**.
  - b. Select the category of performance metric, and then select the performance metric.
  - c. If you want to remove the last metric that you added, click **Remove Metric**.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
12. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

**Related reference:**

“Compare Performance of Multiple Volumes report”

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

### **Compare Performance of Multiple Volumes report**

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

### **Charts**

The charts show the performance metrics for the volumes that you selected for the period that you specified.

## Report output

The following information is provided in the report:

### Performance metrics

Shows the data for each performance metric.

### Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

### Storage volume names

Shows the names of the storage volumes, and the performance metric data for each volume.

### Related tasks:

“Running the Compare Performance of Multiple Volumes report” on page 759  
Use the Compare Performance of Multiple Volumes report to compare up to four performance metrics for multiple volumes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

---

## Predefined reports about storage tiering

Use storage tier reports to help you determine the best storage-tier configuration for your environment.

### Information collection

IBM Tivoli Storage Productivity Center collects information from:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

Tivoli Storage Productivity Center provides Tivoli Common Reporting reports that show performance and usage data for the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified systems in your environment.

### Workload activity reports

Before you run workload activity reports, you must set the back-end properties of the managed disk (MDisk) groups in your environment.

You can use the Tivoli Storage Productivity Center CLI called `tcptool` to set the back-end properties of MDisk groups. Alternatively, you can set the properties in the MDisk Group Details page in the Tivoli Storage Productivity Center stand-alone GUI.

With the information that is collected from Tivoli Storage Productivity Center, you can run workload activity reports. The reports provide information about the workload activity of these resources:

- Managed disk (MDisk) groups
- Virtual disks (VDisks) in MDisk groups
- VDisks in storage resource groups
- VDisks in storage virtualizers

By analyzing workload activity reports, you can determine:

- Whether the workload activity of an MDisk group is too low or too high

- Whether a VDisk should be migrated in one of the following ways:
  - To an MDisk group with a faster read I/O capability rate
  - To an MDisk group with a slower read I/O capability rate

**Explanation:** In the predefined reports about performance, resources, and historical information, MDisk groups are called pools, and VDIsks are called volumes.

**Related information:**

 [setarray](#)

Use the setarray command to set the type of back-end storage system, type of Redundant Array of Independent Disks (RAID), type of disk, and number of disks for an array. For more information about the setarray command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_setarray.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_setarray.html).

 [autosetarray](#)

Use the autosetarray command to extract and save information about storage pools. For more information about the autosetarray command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_autosetarray.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_autosetarray.html).

 [lsarray](#)

Use the lsarray command to display information about arrays and back-end storage subsystems. For more information about the lsarray command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_lsarray.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_lsarray.html).

## Running the Storage Resource Group - VDisk Workload Activity Details report

To see which virtual disks (VDIsks) are most active in a storage resource group, run the Storage Resource Group - VDisk Workload Activity Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **Storage Resource Group - VDisk Workload Activity Details**, then click **Next**.
5. Select a storage virtualizer, and then select a storage resource group.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

**Related reference:**

“Storage Resource Group - VDisk Workload Activity Details report” on page 763 Shows the workload activity of virtual disks (VDIsks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDIsks in the group.



## Storage Resource Group - VDisk Workload Activity Details report

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting period

The name of the storage virtualizer, the name of the storage resource group, the date filter such as Last 7 days, and the reporting period are displayed in the report.

### Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

### VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

### Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

#### Most Active VDisks chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:  
`average(vdisk_peak_utilization)`

To calculate maximum VDisk utilization, the following formula is used:  
`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30

VDisk Report Date	VDisk Peak Utilization (%)
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:  
 $86.30 + 50.01 + 79.41 = 215.72$
2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.  
 $215.72 \div 3 = 71.91$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

#### Least Active VDIsks chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active VDIsks chart are also used in the least active VDIsks chart.

The number of VDIsks that a storage resource group uses determines how many VDIsks are displayed in each chart, as shown in the table.

Number of VDIsks that the storage resource group uses	Number of VDIsks in the Most Active VDIsks chart	Number of VDIsks in the Least Active VDIsks chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

**Tip:** Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of

the MDisk group in the report to open the MDisk Group Details Report. To monitor the workload activity of VDIs for all storage resource groups in a storage virtualizer, run the Storage Resource Group - VDisk Workload Activity Report.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

The following information is provided about VDIs in the storage resource group:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Capacity Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

### **VDisk Read I/O Rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

- VDisk Max Read I/O Rate (ops/s)**  
The maximum read I/O rate of the VDisk in operations per second.
- VDisk Average Read Response Time (ms/op)**  
The average read response time in milliseconds per operation.
- VDisk Max Read Response Time (ms/op)**  
The maximum read response time in milliseconds per operation.
- VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.
- VDisk Write I/O Rate (ops/s)**  
The write I/O rate in operations per second.
- VDisk Average Write Response Time (ms/op)**  
The average write response time in milliseconds per operation.
- VDisk Max Write Response Time (ms/op)**  
The maximum write response time in milliseconds per operation.
- VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.
- VDisk Max Read Density Rate (ops/s per GB)**  
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.
- VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**  
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.
- VDisk Capacity Share of Total VDisk Capacity (%)**  
The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.
- VDisk Average Read Cache Hits (%)**  
The average read cache hits in percent.
- VDisk Max Read Cache Hits (%)**  
The maximum read cache hits in percent.
- VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.
- VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.
- VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.
- VDisk Max Write Cache Delay I/O Rate (ops/s)**  
The maximum write cache delay I/O rate of the VDisk in operations per second.

**Related tasks:**

“Running the Storage Resource Group - VDisk Workload Activity Details report” on page 762  
To see which virtual disks (VDisks) are most active in a storage resource group, run the Storage Resource Group - VDisk Workload Activity Details report.

## Running the VDisk Details report

To see information about a virtual disk (VDisk), run the VDisk Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Details**.
4. Select a storage virtualizer, and then select a VDisk.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

6. Click **Finish**.

### Related reference:

“VDisk Details report”

Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

## VDisk Details report

Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting period

The name of the storage virtualizer, the name of the VDisk, the date filter such as Last 7 days, and the reporting period are displayed in the report.

### Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

### Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

The following information is displayed about the VDisk:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Capacity Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

### **VDisk Read I/O Rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

### **VDisk Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

### **VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

### **VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

### **VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

### **VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

**Related tasks:**

“Running the VDisk Details report” on page 767

To see information about a virtual disk (VDisk), run the VDisk Details report.

## Running the MDisk Groups - VDisk Workload Activity report

To see which virtual disks (VDisks) are most active in managed disk (MDisk) groups, run the MDisk Groups - VDisk Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **MDisk Groups - VDisk Workload Activity**, then click **Next**.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

**Related reference:**

“MDisk Groups - VDisk Workload Activity report”

Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

## MDisk Groups - VDisk Workload Activity report

Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting period

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are displayed in the report.

### Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

### VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

### Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are shown:

#### Most Active MDisk Groups chart

For each MDisk group two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

`average(vdisk_peak_utilization)`



To calculate maximum VDisk utilization, the following formula is used:  
`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:  
 $86.30 + 50.01 + 79.41 = 215.72$
2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.  
 $215.72 \div 3 = 71.91$

In the sample, the average peak VDisk utilization is 71.91.

For each MDisk group, the maximum VDisk utilization is also shown in the chart. The maximum VDisk utilization is the highest value for a VDisk in the VDisk Utilization (%) column. In the sample, the highest value for VDisk1 is 86.30.

#### Least Active MDisk Groups chart

For each MDisk group two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active MDisk group chart are also used in the least active MDisk group chart.

The number of MDisk groups in the storage virtualizer determines how many MDisk groups are shown in each chart, as shown in the following table.

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

**Tip:** Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, click the name of the MDisk group in the report. Click the name of the VDisk in the report to open the VDisk Details Report.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each VDisk in an MDisk group, the following information is provided:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Capacity Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

### **VDisk Read I/O Rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

**VDisk Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

**VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

**Related tasks:**

“Running the MDisk Groups - VDisk Workload Activity report” on page 769  
To see which virtual disks (VDisks) are most active in managed disk (MDisk)

groups, run the MDisk Groups - VDisk Workload Activity report.

## Running the Storage Virtualizer - VDisk Workload Activity report

To see which virtual disks (VDisks) are most active on a storage virtualizer, run the Storage Virtualizer - VDisk Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **Storage Virtualizer - VDisk Workload Activity**, then click **Next**.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

### Related reference:

“Storage Virtualizer - VDisk Workload Activity report”

Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

## Storage Virtualizer - VDisk Workload Activity report

Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting period

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are shown in the report.

### Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

### VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are shown:

### Most Active VDIsks chart

For each storage virtualizer, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

`average(vdisk_peak_utilization)`

To calculate maximum VDisk utilization, the following formula is used:

`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:  
 $86.30 + 50.01 + 79.41 = 215.72$
2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.  
 $215.72 \div 3 = 71.91$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

### Least Active VDIsks chart

For each storage virtualizer, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active VDIsks chart are also used in the least active VDIsks chart.

The number of VDIsks in a storage virtualizer determines how many VDIsks are shown in each chart, as shown in the following table.

Number of VDIsks in the storage virtualizer	Number of VDIsks in the Most Active VDIsks chart	Number of VDIsks in the Least Active VDIsks chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

**Tip:** Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIsks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value}) / 100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each VDisk in a storage virtualizer, the following information is shown:

**MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

**VDisk Report Date**

The date that the data was collected for the report.

**VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

**VDisk Thin Provisioning Capacity Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

**VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

**VDisk Read I/O Rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

**VDisk Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

**VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIsks in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

**Related tasks:**

“Running the Storage Virtualizer - VDisk Workload Activity report” on page 774  
To see which virtual disks (VDIs) are most active on a storage virtualizer, run the Storage Virtualizer - VDisk Workload Activity report.

## Running the Storage Resource Group - VDisk Workload Activity report

To see which virtual disks (VDIs) are most active in storage resource groups, run the Storage Resource Group - VDisk Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **Storage Resource Group - VDisk Workload Activity**.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

**Related reference:**

“Storage Resource Group - VDisk Workload Activity report”

Shows the workload activity of virtual disks (VDIs) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDIs in the group.

## Storage Resource Group - VDisk Workload Activity report

Shows the workload activity of virtual disks (VDIs) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDIs in the group.



**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

## Reporting period

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are shown in the report.

## Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

## VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are shown:

### Most Active Storage Resource Groups chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

`average(vdisk_peak_utilization)`

To calculate maximum VDisk utilization, the following formula is used:

`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:  
 $86.30 + 50.01 + 79.41 = 215.72$

- The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$215.72 \div 3 = 71.91$$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

### Least Active Storage Resource Groups chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active storage resource groups chart are also used in the least active storage resource groups chart.

The number of storage resource groups in a storage virtualizer determines how many storage resource groups are shown in each chart, as shown in the table.

Number of storage resource groups in the storage virtualizer	Number of storage resource groups in the Most Active Storage Resource Groups chart	Number of storage resource groups in the Least Active Storage Resource Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

**Tip:** Click the bar that represents a storage resource group to open the Storage Resource Group: VDisk Workload Activity Details Report. Alternatively, click the name of the storage resource group in the report. Click the name of the VDisk in the report to open the VDisk Details Report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

### Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each storage resource group in a storage virtualizer, the following information is provided about VDIs:

### Storage Resource Group

The name of the storage resource group.

### MDisk Group

The name of the MDisk group.

**VDisk** The name of the VDisk.

### VDisk Report Date

The date that the data was collected for the report.

### VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

### VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

### VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

### VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

### VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

### VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

**VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

**Related tasks:**

“Running the Storage Resource Group - VDisk Workload Activity report” on page 778

To see which virtual disks (VDIs) are most active in storage resource groups, run the Storage Resource Group - VDisk Workload Activity report.

## Running the MDisk Group - VDisk Workload Activity Details report

To see which of the virtual disks in a managed disk (MDisk) group are the most active, run the MDisk Group - VDisk Workload Activity Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **MDisk Group - VDisk Workload Activity Details**.
4. Select a storage virtualizer, and then select an MDisk group.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

6. Click **Finish**.

### Related reference:

“MDisk Group - VDisk Workload Activity Details report”

Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

## MDisk Group - VDisk Workload Activity Details report

Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting period

The name of the storage virtualizer, the name of the MDisk group, the date filter such as Last 7 days, and the reporting period are displayed on the report.

### Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

### VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

### Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For

example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are shown:

#### Most Active VDisks chart

For each MDisk group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

`average(vdisk_peak_utilization)`

To calculate maximum VDisk utilization, the following formula is used:

`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:  
 $86.30 + 50.01 + 79.41 = 215.72$
2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.  
 $215.72 \div 3 = 71.91$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

#### Least Active VDisks chart

For each MDisk group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active MDisk group chart are also used in the least active MDisk group chart.

The number of VDisks in an MDisk group determines how many VDisks are shown in each chart, as shown in the following table.

Number of VDIs in the MDisk group	Number of VDIs in the Most Active VDIs chart	Number of VDIs in the Least Active VDIs chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

**Tip:** Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each VDisk in an MDisk group, the following information is provided:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

**VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

**VDisk Thin Provisioning Capacity Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

**VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

**VDisk Read I/O Rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

**VDisk Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

**VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.



**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

**Related tasks:**

“Running the MDisk Group - VDisk Workload Activity Details report” on page 783

To see which of the virtual disks in a managed disk (MDisk) group are the most active, run the MDisk Group - VDisk Workload Activity Details report.

## Running the MDisk Group Details report

To see statistics about the performance of a managed disk (MDisk) group, run the MDisk Group Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **MDisk Group Details**.
4. Select a storage virtualizer, and then select an MDisk group.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

6. Click **Finish**.

**Related reference:**

“MDisk Group Details report”

Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

## MDisk Group Details report

Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting period

The name of the storage virtualizer, the name of the MDisk group, the date filter such as Last 7 days, and the reporting period are displayed in the report.

## Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

## Read I/O rate capability

If you do not set the back-end pool properties of an MDisk group:

- Not Set is displayed in the MDisk Group Read I/O Rate Capability column
- Not Available is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

## Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value}) \div 100$ .

**Tip:** To monitor the workload activity of all MDisk groups in a storage virtualizer, run the MDisk Groups - Workload Activity report.

## Report output

The following information is provided about the MDisk group:

### MDisk Group

The name of the MDisk group.

### MDisk Group Report Date

The date that the data was collected for the report.

### MDisk Group Capacity Status

If the storage capacity that is allocated is less than the actual storage capacity, a status value of OK is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of Warning is displayed. The actual capacity of the MDisk group is calculated as follows:

$(\text{mdisk\_group\_real\_capacity\_value} \div \text{mdisk\_group\_capacity\_value}) \times 100$

**MDisk Group Capacity (GB)**

The total storage capacity in GB of all of the disks in the MDisk group.

**MDisk Group Real Capacity (GB)**

The total amount in GB of real storage capacity in the MDisk group.

**MDisk Group Free Space (GB)**

The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

**MDisk Group Available Size < Actual Size Threshold of {%**

If the amount of storage space that is available for allocation to VDIs is less than the actual size threshold, Yes is displayed in the column. If the amount of storage space that is available for allocation to VDIs is more than the actual size threshold, No is displayed in the column.

**MDisk Group Read I/O Rate (ops/s)**

The read I/O rate of the MDisk group in operations per second.

**MDisk Group Read I/O Rate Capability (ops/s)**

The read I/O capability rate of the MDisk group in operations per second.

**Tip:** To find out how the read I/O capability rate is calculated, see “Read I/O capability formula” on page 795

**MDisk Group Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the MDisk group in operations per second.

**MDisk Group Read I/O Rate Availability (ops/s)**

The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

**MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {%**

The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, Yes is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, No is displayed in the column.

**MDisk Group Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**MDisk Group Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**MDisk Group Max Write I/O Rate (ops/s)**

The maximum write I/O rate in operations per second.

**MDisk Group Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**MDisk Group Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**MDisk Group Average Write I/O Size (KB)**

The average size of write I/O operations in KB.

**Related tasks:**

“Running the MDisk Group Details report” on page 787

To see statistics about the performance of a managed disk (MDisk) group, run the MDisk Group Details report.

## Running the MDisk Groups - Workload Activity report

To see which managed disk (MDisk) groups are most active, run the MDisk Groups - Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **MDisk Groups - Workload Activity**.
4. Select a storage virtualizer.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

**Tip:** To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

6. Click **Finish**.

**Related reference:**

“MDisk Groups - Workload Activity report”

Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

“Read I/O capability formula” on page 795

The read I/O capability formula is used by IBM Tivoli Storage Productivity Center to estimate the workload capacity of managed disk (MDisk) groups.

## MDisk Groups - Workload Activity report

Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

**Restriction:** This report applies only to SAN Volume Controller and Storwize systems.

### Reporting periods

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are displayed in the report.

### Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

## Read I/O rate capability

To estimate the read I/O rate capability of an MDisk group, you must set the back-end pool properties of MDisk groups.

If you do not set the back-end pool properties of an MDisk group:

- Not Set is displayed in the MDisk Group Read I/O Rate Capability column
- Not Available is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

- Most Active MDisk Groups
- Most Available MDisk Groups

### Most Active MDisk Groups

For each MDisk group two bars are used:

- A yellow bar that shows the average activity value
- A red bar that shows the maximum activity value

To calculate the average activity value of an MDisk group, the following formula is used:

$$\text{average}(\text{maximum\_read\_I/O\_rate} \div \text{read\_I/O\_rate\_capability})$$

To calculate the maximum activity value, the following formula is used:

$$\text{maximum}(\text{maximum\_read\_I/O\_rate} \div \text{read\_I/O\_rate\_capability})$$

In this sample excerpt of report output for mdiskgroup1, the values that are required to calculate average and maximum activity values are provided.

MDisk Group Report Date	MDisk Group Read I/O Rate Capability (ops/s)	MDisk Group Max Read I/O Rate (ops/s)
Dec 6 2011	331	216.77
Dec 7 2011	331	256.40
Dec 8 2011	331	270.87

The average activity value of an MDisk group is calculated as follows:

1. For each report in the MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:

$$216.77 \div 331 = 0.65$$

$$256.40 \div 331 = 0.77$$

$$270.87 \div 331 = 0.81$$

2. The result of the operation for each report is added:

$$0.65 + 0.77 + 0.81 = 2.23$$

- To get the average value, the total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$2.23 \div 3 = 0.74$$

The maximum value of an MDisk group is calculated as follows. For each report in an MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:

$$216.77 \div 331 = 0.65$$

$$256.40 \div 331 = 0.77$$

$$270.87 \div 331 = 0.81$$

In this sample, the maximum value is 0.81.

### Most Available MDisk Groups

For each MDisk group two bars are used:

- A blue bar to represent minimum values
- A red bar to represent average values

In this sample excerpt of report output for mdiskgroup2, the values that are required to calculate minimum and average availability values for MDisk Groups are provided.

MDisk Group Report Date	MDisk Group Read I/O Rate Availability (ops/s)
Dec 6 2011	435.41
Dec 7 2011	560.92
Dec 8 2011	480.95

The minimum value of an MDisk group is calculated as follows:

- For all of the reports for an MDisk group, the lowest value in the MDisk Group Read I/O Rate Availability column is used. In this sample, the lowest value for mdiskgroup2 is 435.41.

The average value of an MDisk group is calculated as follows:

- For each report in an MDisk group, the values in the MDisk Group Read I/O Rate Availability column are added:

$$435.41 + 560.92 + 480.95 = 1477.28$$

- To get the average value, the total is divided by the number of days that are specified in the date range. (For example, if the reporting period spans 3 days, the total is divided by 3.)

$$1477.28 \div 3 = 492.42$$

In this sample, the average value is 492.42.

The number of MDisk groups in a storage virtualizer determines how many MDisk groups are displayed in each chart, as shown in the following table.

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
≥ 10	5	5
9	5	4

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

**Tip:** Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, you can click the name of the MDisk group in the report.

## Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$$(maximum\_response\_time\_value \times maximum\_response\_time\_threshold\_value) \div 100.$$

The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each MDisk group, the following information is provided:

### MDisk Group

The name of the MDisk group.

### MDisk Group Report Date

The date that the data was collected for the report.

### MDisk Group Capacity Status

If the storage capacity that is allocated is less than the actual storage capacity, a status value of OK is displayed. If the storage capacity that is

allocated is more than the actual storage capacity, a status value of Warning is displayed. The actual capacity of the MDisk group is calculated as follows:

$$(\text{mdisk\_group\_real\_capacity\_value} \div \text{mdisk\_group\_capacity\_value}) \times 100$$

**MDisk Group Capacity (GB)**

The total storage capacity in GB of all of the disks in the MDisk group.

**MDisk Group Real Capacity (GB)**

The total amount in GB of real storage capacity in the MDisk group.

**MDisk Group Free Space (GB)**

The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

**MDisk Group Available Size < Actual Size Threshold of {%}**

If the amount of storage space that is available for allocation to VDIsks is less than the actual size threshold, Yes is displayed in the column. If the amount of storage space that is available for allocation to VDIsks is more than the actual size threshold, No is displayed in the column.

**MDisk Group Read I/O Rate (ops/s)**

The read I/O rate of the MDisk group in operations per second.

**MDisk Group Read I/O Rate Capability (ops/s)**

The read I/O capability rate of the MDisk group in operations per second.

**Tip:** To find out how the read I/O capability rate is calculated, see “Read I/O capability formula” on page 795

**MDisk Group Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the MDisk group in operations per second.

**MDisk Group Read I/O Rate Availability (ops/s)**

The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

**MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {%}**

The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, Yes is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, No is displayed in the column.

**MDisk Group Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**MDisk Group Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**MDisk Group Max Write I/O Rate (ops/s)**

The maximum write I/O rate in operations per second.

**MDisk Group Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.



**MDisk Group Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**MDisk Group Average Write I/O Size (KB)**

The average size of write I/O operations in KB.

**Related tasks:**

“Running the MDisk Groups - Workload Activity report” on page 790

To see which managed disk (MDisk) groups are most active, run the MDisk Groups - Workload Activity report.

## Read I/O capability formula

The read I/O capability formula is used by IBM Tivoli Storage Productivity Center to estimate the workload capacity of managed disk (MDisk) groups.

The estimated workload capacity is compared with the actual workload capacity to assess the most active and least active MDisk groups in a storage subsystem. To monitor the workload activity of MDisk groups, you can run customized reports such as the MDisk Groups - Workload Activity Report. This formula is used for MDisk groups in the following storage systems:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

The formula that is used by Tivoli Storage Productivity Center to estimate the read I/O capability rate is:

$$\frac{[\text{Estimated Read I/O Workload} * \text{Type of Disk I/O Capability} * \text{Number of Disks}]}{[\text{Estimated Read I/O Workload} * (1 - \text{Cache Hit Ratio}) + \text{Estimated Write I/O Workload} * \text{RAID Type Weighted I/O Ratio}]}$$

There are two types of value that are used by the formula:

- Fixed values provided by Tivoli Storage Productivity Center
- Variable values that you must specify for each MDisk group

The following fixed values are used by the formula:

**Estimated Read I/O Workload**

The estimated read I/O workload of an MDisk group. The estimate is set to a fixed value of 0.7 (70%).

**Estimated Write I/O Workload**

The estimated write I/O workload of an MDisk group. The estimate is set to a fixed value of 0.3 (30%).

You specify the following variable values for each MDisk group:

**Type of Disk I/O Capability**

When you specify the type of disk that the back-end pool uses:

- You can accept the default I/O capability per second of the disk
- You can enter a new value for the I/O capability per second of the disk

For example, the default value for the I/O capability rate per second of fiber types of disk with 15,000 revolutions per minute (rpm) is 150 I/O operations per second. To change the I/O capability for a type of disk, you issue the **setbackdisktype** command and you enter a new value for the **-iops** parameter.

### Number of Disks

You specify the number of physical disks that the back-end pool uses.

### Cache Hit Ratio

When you specify the type of storage system that the back-end pool uses:

- You can accept the default cache hit ratio for the storage system
- You can enter a new value for the cache hit ratio of the storage system

For example, the default cache hit ratio of the XIV system is 0.7 (70%). To change the cache hit ratio for a type of storage system, you issue the **setbackendtype** command and you enter a new value for the **-cachehit** parameter.

### RAID Type Weighted I/O Ratio

When you specify the type of RAID that the back-end pool uses:

- You can accept the default value for the weighted I/O ratio of the type of RAID
- You can enter a new value for the weighted I/O ratio of the type of RAID

For example the weighted I/O ratio for RAID5 types of RAID is 4:1. To change the weighted I/O ratio for a type of RAID, you issue the **setbackraidtype** command and you enter a new value for the **-weightedio** parameter.

**Tip:** In the stand-alone GUI, you can set the values for the back-end properties of an MDisk group. You can also use the Tivoli Storage Productivity Center CLI called **tcptool** to set the values for back-end MDisk groups. See the **setarray** command. To change the variable values that are used by the Read I/O capability formula, you must use the Tivoli Storage Productivity Center CLI.

### Example

The following values have been set for a back-end pool. The type of storage system is the XIV system. The type of disk is F15 (Fiber - 15 000 rpm), and the type of RAID is X.

Name of variable	Value
Cache Hit Ratio	0.5 (50%)
Type of Disk I/O Capability	150
RAID Type Weighted I/O Ratio	4
Number of Disks	10

Based on the values that are provided, the read I/O capability per second is:

$$[0.7 * 150 * 10] / [0.7 * (1 - 0.5) + 0.3 * 4]$$

The read I/O capability per second rate is 677. (The result of the calculation is rounded to the nearest whole number.)

**Related information:**

 [setarray](#)

Use the `setarray` command to set the type of back-end storage system, type of Redundant Array of Independent Disks (RAID), type of disk, and number of disks for an array. For more information about the `setarray` command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_setarray.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_setarray.html).

 [setbackendtype](#)

Use the `setbackendtype` command to set or update the type of back-end storage system. For more information about the `setbackendtype` command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_setbackendtype.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_setbackendtype.html).

 [setbackdisktype](#)

Use the `setbackdisktype` command to set or update the type of back-end disk. For more information about the `setbackdisktype` command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_setbackdisktype.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_setbackdisktype.html).

 [setbackraidtype](#)

Use the `setbackraidtype` command to set the types of back-end RAID arrays for managed disk groups. For more information about the `setbackraidtype` command, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_r\\_cli\\_setbackraidtype.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_r_cli_setbackraidtype.html).

## Changing threshold values in storage tier reports

Use the `setdscfg` command to change the threshold values in workload activity reports for managed disk groups and virtual disks.

To change threshold values, you must know the value of the property key that is associated with the name of the threshold in the report.

1. To start the Tivoli Storage Productivity Center CLI, issue the `tpctool` command.
2. Issue the `setdscfg` command, and then set the following values by using the `-property` parameter:
  - a. Specify the name of the threshold with the `property_key` variable.
  - b. Specify the new threshold value with the `value` variable.

The following sample code sets the read I/O rate for MDisk groups to 80%:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd ***** -property  
ManagedDiskGroupAvailIORate 80
```

3. Press Enter.

The value of the threshold is changed.

**Related reference:**

“Threshold values in storage tier reports” on page 798

Threshold values are used to assess the workload activity of managed disk (MDisk) groups and virtual disks (VDisks) in storage tier reports. To change a threshold value, you must know the value of the property key that corresponds with the threshold value that is used in the report.

## Threshold values in storage tier reports

Threshold values are used to assess the workload activity of managed disk (MDisk) groups and virtual disks (VDisks) in storage tier reports. To change a threshold value, you must know the value of the property key that corresponds with the threshold value that is used in the report.

The following table contains the name of the threshold value and the corresponding value of the property key for the following reports:

- The MDisk Groups - Workload Activity Report
- The MDisk Group Details Report

Threshold name in report	Property key value
Actual Size Threshold	<i>ManagedDiskGroupAvailSize</i>
Read I/O Rate Threshold	<i>ManagedDiskGroupAvailIORate</i>
Max Read Response Time Threshold	<i>ManagedDiskGroupBERResponseExceeds</i>
Max Write Response Time Threshold	<i>ManagedDiskGroupBEWResponseExceeds</i>

The following table contains the name of the threshold value and the corresponding value of the property key for the following reports:

- MDisk Groups - VDisk Workload Activity Report
- MDisk Groups - VDisk Workload Activity Details Report
- VDisk Details Report
- Storage Resource Group - VDisk Workload Activity Report

Threshold name in report	Property key value
VDisk Max Read Response Time Threshold	<i>VDiskFERResponseExceeds</i>
VDisk Max Write Response Time Threshold	<i>VDiskFEWResponseExceeds</i>
Max Read Cache Hits Time Threshold	<i>VDiskFERHitExceeds</i>
Max Read I/O Rate Threshold	<i>VDiskFEIORExceeds</i>
Max Write I/O Rate Threshold	<i>VDiskFEIOWExceeds</i>

### Related tasks:

“Changing threshold values in storage tier reports” on page 797

Use the **setdscfg** command to change the threshold values in workload activity reports for managed disk groups and virtual disks.

---

## Custom reports about performance

You can use the IBM Cognos reporting software to create custom performance reports. Performance reports can contain detailed information about the performance of monitored resources, and some of the properties of those resources.

The resources for which you can create performance reports are storage systems and their components, storage virtualizers and their components, and switches and switch ports. You can add statistics for volumes, back-end arrays, caches, ports, port errors, and other statistics to performance reports. For example, you can add key statistics to a report to create and generate a performance chart for a specific volume of a storage system.

All the performance metrics from the stand-alone GUI are available in the new reporting interface.

## Creating custom performance reports

You can create reports that show performance metrics. Performance reports can contain detailed information about the performance of monitored resources and some of the properties of those resources. You access custom performance reports from the Tivoli Storage Productivity Center web-based GUI, and you create the reports in the IBM Cognos reporting software.


When you create a performance report, include properties and metrics for only one type of device in the report.

The list that you use to create performance reports does not include information about the relationships between storage resources. If you include different types of resources in the report, the report might return incorrect results. For example, if you include storage virtualizers and storage systems in a report, the report might return incorrect results.

**Tip:** Before you run a report, you can preview the report with no data.


You must specify the time interval for which you want to view performance information. You can specify the time interval in the following ways:

- Select a time unit, for example, day, day of week, or week of month.
- Select performance metrics from sample data, hourly data, or daily data in the resource for which you want performance information. You can also add the time and the interval length to reports. The time is the time at which the performance data was collected. The interval length is the length of the data interval in seconds for sample data, or in minutes for hourly data or daily data.

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the type of report that you want to see, select one of the following items:
  - To create a basic report with limited formatting about the performance of resources, select **Cognos > Create performance reports**.
  - To create a report about the performance of resources and apply advanced formatting features to the report, complete the following steps:
    - a. Select **Cognos > Create performance reports**.
    - b. Click the **Return** icon .
    - c. Click **Launch > Report Studio**.
3. Expand **Tivoli Storage Productivity Center** in the list.
4. Expand **Performance**.
5. Optional: Select a time unit from the **Time** item.
6. Explore the folders in the list to locate the resources, properties, and metrics that you want to include in the report.
7. Depending on your browser, add properties and metrics to the report:
  - In Internet Explorer, drag the items that you want to see information about to the work area.
  - In Firefox, select the items that you want to see information about, and then click **Insert**.

8. Use filters with your performance reports to ensure that your reports do not include too much data. To use filters with your performance reports, do the following steps:
  - a. Add a resource or resource type to the report.
  - b. Select the column heading for the resource or resource type in the work area, and then filter the name of the resource or resource type if necessary.
  - c. Locate the Sample Data, Hourly Data, or Daily Data folder for the resource or resource type in the list. Add a time stamp from the list to the work area. For example, you can add a day time stamp to the report.
  - d. Select the column heading for the time stamp in the work area, and then filter the time. If you do not enter a time stamp, the report is run on all the records for this resource or resource type in Tivoli Storage Productivity Center.
9. Click **Save**.

**Related information:**

 **Run a Report**

For more information about how to preview and run reports, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cr\\_rptstd.10.2.0.doc/t\\_cr\\_rptstd\\_intrd\\_run\\_report.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cr_rptstd.10.2.0.doc/t_cr_rptstd_intrd_run_report.html).

## Data and properties in performance reports

You can include capacity data, properties, and other information about storage systems, storage virtualizers, and switches in performance reports.

### Data for storage systems in performance reports

You can include general information, capacity data, properties, and other information about storage systems in performance reports.

### Information about storage systems

You can create performance reports that include the following information:

#### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Capacity and usage data

You can create performance reports that include the following information:

#### Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

#### Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

#### Storage System Available Pool Space (GiB)

The amount of unused space in pools that are on the storage system.

#### Storage System Volume Capacity (GiB)

The amount of space on all volumes on the storage system.

**Storage System Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage System Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage System Unallocated Disk Space (GiB)**

The amount of disk space that can be added to a pool.

**Storage System Number of Disks**

The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

**Storage System Number of Volumes**

The number of volumes on a resource.

**Storage System Cache (GiB)**

The size of the cache on the storage system. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Real Available Pool Space (GiB)**

The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Real Configured Pool Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Volume Capacity for z/OS (GiB)**

The amount of space on all volumes on the storage system that the z/OS operating system can use.

**Storage System Volume Capacity Assigned to MDisks (GiB)**

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

**Storage System Disk Capacity (GiB)**

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Total Disk Capacity (GiB)**

The amount of space on physical disks on a storage system, including spare disks.

**Storage System Physical Allocation Percentage**

The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated pool space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

### **Storage System Virtual Allocation Percentage**

The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

### **Storage System Shortfall Percentage**

The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB  $\div$  100 GiB  $\times$  100).

This value is only available for pools with thin-provisioned volumes.

### **Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### **Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.



## Component properties

You can create performance reports that include the following information:

### Storage System Access Information

The URL from which you can access the storage system or storage virtualizer.

### Storage System Alias

An alternative name for a storage system. The user defines this name.

### Storage System Code Level

The Shared Ethernet Adapter level of a DS8000 storage system. For other storage systems, this value is the firmware version.

### Storage System Time Zone

The time zone in which a resource is located.

### Storage System Custom Tag 1, 2, and 3

User-defined text that is associated with a storage system. You can add or edit the custom tags for a storage system in the **Properties** pane of the storage system in the Tivoli Storage Productivity Center web-based GUI.

### Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

### Storage System User Provided Name

The name that was specified for the storage system in the storage environment.

### Storage System Is Compression Active

Shows whether the compression feature is enabled on the storage system. If this value is Yes, the compression feature is enabled.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

## Status information

You can create performance reports that include the following information:

### Storage System Status

The condition of the resource, for example normal, warning, or error.

### Storage System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### Storage System Last Data Collection

The date and time when storage statistics were last collected from the resource.

**Storage System Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

**Vendor, model, and device information**

You can create performance reports that include the following information:

**Storage System Type**

The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

**Storage System Vendor**

The vendor who supplied the resource.

**Storage System Model**

The model name or model number of the storage system.

**Storage System Machine Type**

The machine type of a storage system. For example, the storage system can be **DS8800**, **DS8700**, **XIV - 2812**, or another type of storage system.

**Storage System IP Address**

The IP address of the resource.

**Storage System Serial Number**

The serial number of the resource.

**Performance metrics for SAN Volume Controller and Storwize systems:**

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for these systems in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

**Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

**Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

**Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

**Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

**Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

**Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

**Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Local Node Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

**Overall Port to Local Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Remote Node Response Time (ms/op)**

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

**Overall Port to Remote Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

**Port error data**

You can create performance reports that include the following information:

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.



**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

## Other data

You can create performance reports that include the following information:

### **Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

### **Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

### **Peak Read Response Time (ms)**

The longest time to complete a read operation.

### **Peak Write Response Time (ms)**

The longest time to complete a write operation.

### **Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

### **Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

### **Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

### **Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

### **Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

### **Zero Buffer-to-Buffer Credit Percentage**

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

### **Performance metrics for DS8000 and DS6000 storage systems:**

You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for DS8000 and DS6000 storage systems in performance reports.

#### **Port data**

You can create performance reports that include the following information:

##### **Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

##### **Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

##### **Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

##### **Port PPRC Send I/O Rate (ops/s)**

The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.

##### **Port PPRC Receive I/O Rate (ops/s)**

The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.

##### **Total Port PPRC I/O Rate (ops/s)**

The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.

##### **Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

##### **Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

##### **Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

##### **Port PPRC Send Data Rate (MiB/s)**

The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.

##### **Port PPRC Receive Data Rate (MiB/s)**

The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.

**Total Port PPRC Data Rate (MiB/s)**

The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.

**Port Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation.

**Port Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation.

**Overall Port Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation.

**Port PPRC Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.

**Port PPRC Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.

**Overall Port PPRC Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

**Port Send Transfer Size (KiB/op)**

The average number of KiB that are transferred per send operation.

**Port Receive Transfer Size (KiB/op)**

The average number of KiB that are transferred per receive operation.

**Overall Port Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

**Port error data**

You can create performance reports that include the following information:

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases,

this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Error Frame Rate (count/s)**

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

**Link Reset Transmitted Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

**Link Reset Received Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

**Out of Order Data Rate (count/s)**

The average number of times per second that an out-of-order frame is detected.

**Out of Order ACK Rate (count/s)**

The average number of times per second that an out-of-order acknowledge (ACK) frame is detected.

An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.

**Duplicate Frame Rate (count/s)**

The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.

**Invalid Relative Offset Rate (count/s)**

The average number of times per second that frames are received with an invalid relative offset in the frame header.

**Sequence Timeout Rate (count/s)**

The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.

**Credit Recovery Link Reset Rate**

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

## Volume data

You can create performance reports that include the following information:

### **Normal Read I/O Rate (ops/s)**

The average number of nonsequential read operations per second.

### **Sequential Read I/O Rate (ops/s)**

The average number of sequential read operations per second.

### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

### **Normal Write I/O Rate (ops/s)**

The average number of nonsequential write operations per second.

### **Sequential Write I/O Rate (ops/s)**

The average number of sequential write operations per second.

### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

### **Total Normal I/O Rate (ops/s)**

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

### **Total Sequential I/O Rate (ops/s)**

The average number of sequential I/O operations per second. This value includes both read and write operations.

### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

### **Record Mode Read I/O Rate (ops/s)**

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Normal Read Cache Hit Percentage**

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Sequential Read Cache Hit Percentage**

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Normal Write Cache Hit Percentage**

The percentage of nonsequential write operations that are handled in the cache.

**Sequential Write Cache Hit Percentage**

The percentage of sequential write operations that are handled in the cache.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Normal Cache Hit Percentage**

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.



You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Sequential Cache Hit Percentage**

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Record Mode Read Cache Hit Percentage**

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

**Other data**

You can create performance reports that include the following information:

**HPF Read I/O Rate (ops/s)**

The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

**HPF Write I/O Rate (ops/s)**

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

**Total HPF I/O Rate (ops/s)**

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

**HPF I/O Percentage**

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

**Cache Holding Time (s)**

The average number of seconds that I/O data for a storage system controller is held in the cache. A short cache-holding time indicates adverse performance.

**PPRC Transfer Rate (ops/s)**

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

**Performance metrics for an XIV system:**

You can include performance data for volumes, caches, and other data for an XIV system in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

## Cache data

You can create performance reports that include the following information:

### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

### Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### Read Data Cache Hit Percentage

The percentage of all read data that is read from the cache.

### Write Data Cache Hit Percentage

The percentage of all write data that is written to cache slots that are marked as modified.

### Overall Data Cache Hit Percentage

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

### SSD Read Cache Hit Percentage

The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

### SSD Read Data Cache Hit Percentage

The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

### Read Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a read-cache hit operation.

### Write Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a write-cache hit operation.

**Overall Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

**Read Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a read-cache miss operation.

**Write Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a write-cache miss operation.

**Overall Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

**SSD Read Cache Hit Response Time (ms/op)**

The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**Other data**

You can create performance reports that include the following information:

**Small Transfers I/O Percentage**

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers I/O Percentage**

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers I/O Percentage**

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers I/O Percentage**

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

**Performance metrics for block server systems:**

You can include performance metrics for volumes, and cache data for block server storage systems in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Data for storage system volumes in performance reports**

You can include general information, capacity data, properties, and other information about storage system volumes in performance reports.

**Information about storage system volumes**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Controller Name**

The name that was assigned to the controller when it was added to the system.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage RAID Array Name**

The name that was assigned to the RAID array when it was added to the system.

**Storage I/O Group Name**

The name that was assigned to the I/O group when it was added to the system.

**Storage Preferred Node Name**

The name of the storage node that is used for I/O operations for a volume.

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Storage Volume Assigned Host Connection**

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

**Storage Volume Host Connection OS Type**

The operating system type of the server or hypervisor that the volume is assigned to.

**Storage System Volume Group ID**

The unique identifier for a collection of volumes on DS8000 and DS6000 storage systems. The storage system generates this identifier.

**Storage Volume Group Name**

The user-defined name for a collection of volumes on DS8000 and DS6000 storage systems.

**Server Name**

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

**Hypervisor Name**

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

**Hypervisor Cluster Name**

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage Volume Capacity (GiB)**

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

**Storage Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin

provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Volume Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine this value:

capacity - allocated space

Available only for thin provisioned volumes.

**Storage Volume Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$(\text{allocated space} \div \text{capacity}) \times 100$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Volume Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$(\text{unallocatable volume Space} \div \text{volume unallocated space}) \times 100$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.



This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$$\frac{100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space}}{100 \text{ GiB unallocated volume space}} \times 100$$

#### **Storage Volume Unallocatable Space (GiB)**

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

#### **Storage Volume Used Allocated Space Percentage**

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 volumes.

#### **Storage Volume Unused Space (GiB)**

The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, Storwize V7000, and Storwize V3700 pools.

#### **Storage Volume Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

### **Component properties**

You can create performance reports that include the following information:

#### **Storage Volume RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

#### **Storage Volume Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Volume Is Assigned**

Shows whether the volume is assigned to a server. If this value is Yes, the volume is assigned to a server.

**Storage Volume Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

**Storage Volume Is Encryptable**

Shows whether the resource can be encrypted. If this value is Yes, the resource can be encrypted.

**Storage Volume Is Volume Control Manager Database**

Shows whether a volume is used as the Volume Control Manager Database. If this value is Yes, the volume is used as the Volume Control Manager Database.

**Storage Volume Format**

Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

**Storage Volume Serial Number**

The serial number of the resource.

**Storage Volume Number**

The number of the volume that is assigned by the system.

**Storage Volume WWN**

The worldwide name of the volume.

**Storage Volume FlashCopy Relationship**

Shows whether a volume on a storage system is in a FlashCopy relationship. This property can contain the following values:

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Not in a FlashCopy Relationship**

The volume is not in a FlashCopy relationship.

**Storage Volume Remote Copy Relationship**

Shows whether a volume on a storage system is in a remote copy relationship. This property can contain the following values:

**Not in a Copy Relationship**

The volume is not in a remote copy relationship.

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Storage Volume Logical Subsystem**

The logical subsystem (LSS) to which a volume or pool belongs.

**Storage Volume Type**

The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

**Storage Volume Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Is Auto Expand Enabled**

Shows whether a volume can automatically allocate new extents from a pool. Volumes might allocate new extents to expand the real capacity of the volume. If this value is Yes, the volume can automatically allocate new extents from a pool.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Metro Mirror Name**

The name of the Metro Mirror that keeps the synchronous copy of the storage volume. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Mirror Count**

The number of mirrors that keep a synchronous copy of the resource. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Throttle**

The maximum number of commands that the volume can queue. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Warning Level**

The percentage of volume capacity that is used at which a warning is generated. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Is Compressed**

Shows whether the volume is compressed. If this value is Yes, the volume is compressed.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

## Status information

You can create performance reports that include the following information:

### Storage Volume Status

The condition of the resource, for example normal, warning, or error.

### Storage Volume Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### Storage Volume Native Status

The condition of the volume. The status can be online, offline, degraded, or excluded.

## Performance metrics for volumes on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes, caches, and other data for volumes on these systems in performance reports.

### Volume data

You can create performance reports that include the following information:

#### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

#### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### Maximum Read I/O Rate (ops/s)

The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

#### Maximum Write I/O Rate (ops/s)

The maximum write I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

#### Maximum Write I/O Size (KiB)

The maximum size of write I/O operations in kibibytes. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

#### Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations.

#### Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Maximum Read Response Time (ms/op)**

The maximum read response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Write Response Time (ms/op)**

The maximum write response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Maximum Volume Utilization Percentage**

The maximum period in percentage that the volume is busy. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Maximum Write Cache Delay I/O Rate (ops/s)**

The maximum number of delayed I/O operations per second for the volume, where the delay occurred because of space constraints or other conditions.

This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Maximum Read Cache Hit Percentage**

The maximum number of cache hits in percentage for sequential and nonsequential read operations for the volume over a specified interval. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

### **Nonpreferred Node Usage Percentage**

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

### **Performance metrics for volumes on DS8000 and DS6000 storage systems:**

You can include performance data for volumes, caches, and other data for storage system volumes on DS8000 and DS6000 storage systems in performance reports.

#### **Volume data**

You can create performance reports that include the following information:

#### **Normal Read I/O Rate (ops/s)**

The average number of nonsequential read operations per second.

#### **Sequential Read I/O Rate (ops/s)**

The average number of sequential read operations per second.

#### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

#### **Normal Write I/O Rate (ops/s)**

The average number of nonsequential write operations per second.

#### **Sequential Write I/O Rate (ops/s)**

The average number of sequential write operations per second.

#### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

#### **Total Normal I/O Rate (ops/s)**

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

#### **Total Sequential I/O Rate (ops/s)**

The average number of sequential I/O operations per second. This value includes both read and write operations.

#### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### **Record Mode Read I/O Rate (ops/s)**

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

#### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

#### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.



**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Normal Read Cache Hit Percentage**

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Sequential Read Cache Hit Percentage**

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Normal Write Cache Hit Percentage**

The percentage of nonsequential write operations that are handled in the cache.

**Sequential Write Cache Hit Percentage**

The percentage of sequential write operations that are handled in the cache.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Normal Cache Hit Percentage**

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Sequential Cache Hit Percentage**

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Record Mode Read Cache Hit Percentage**

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

**Other data**

You can create performance reports that include the following information:

**HPF Read I/O Rate (ops/s)**

The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

**HPF Write I/O Rate (ops/s)**

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

**Total HPF I/O Rate (ops/s)**

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

**HPF I/O Percentage**

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

**PPRC Transfer Rate (ops/s)**

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

**Performance metrics for storage system volumes on an XIV system:**

You can include performance data for volumes, caches, and other data for storage system volumes on an XIV system in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Read Data Cache Hit Percentage**

The percentage of all read data that is read from the cache.

**Write Data Cache Hit Percentage**

The percentage of all write data that is written to cache slots that are marked as modified.

**Overall Data Cache Hit Percentage**

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

**Read Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a read-cache miss operation.

**Write Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a write-cache miss operation.

**Overall Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

**Read Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a read-cache hit operation.

**Write Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a write-cache hit operation.

**Overall Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

**SSD Read Cache Hit Percentage**

The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**SSD Read Data Cache Hit Percentage**

The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**SSD Read Cache Hit Response Time (ms/op)**

The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**Other data**

You can create performance reports that include the following information:

**Small Transfers I/O Percentage**

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers I/O Percentage**

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers I/O Percentage**

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers I/O Percentage**

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

**Performance metrics for storage system volumes on block servers:**

You can include performance data for volumes and caches for storage system volumes on block servers in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Volume Utilization**

The average percentage of time that the volume is busy.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Data for storage system ports in performance reports**

You can include general information, properties, and other information about storage system ports in performance reports.

**Information about storage system ports**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Module Name**

The name that was assigned to the module when it was added to the system.

**Storage Port Name**

The name that was assigned to the storage port when the storage system was added to the system.

**Component properties**

You can create performance reports that include the following information:

**Storage System Port Number**

The port number on the resource.

**Storage System Port Type**

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N\_Port, F\_Port, or another type of port.

**Storage System Port Speed (GiB/s)**

The speed of a port, which is measured in GiB per second.

**Storage System Port Location**

The physical location of a port on DS8000, DS6000, TotalStorage Enterprise Storage Server storage systems, and on XIV systems. For all other systems, this property shows the WWPN of the port.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Status information**

You can create performance reports that include the following information:

**Storage System Port Status**

The condition of the resource, for example normal, warning, or error.

**Storage System Port Enabled State**

Shows whether a port is enabled, disabled, or is enabled but offline.

**Storage System Port Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Performance metrics for ports on SAN Volume Controller and Storwize systems:**

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports and port errors on these systems in performance reports.



## Port data

You can create performance reports that include the following information:

### **Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

### **Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

### **Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

### **Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

### **Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

### **Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

### **Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

### **Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

### **Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

### **Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

### **Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

### **Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

### **Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

### **Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port error data**

You can create performance reports that include the following information:

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

**Zero Buffer-to-Buffer Credit Percentage**

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

**Port Congestion Index**

The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

**Performance metrics for ports on DS8000 and DS6000 storage systems:**

You can include performance data for storage system port and port errors on DS8000 and DS6000 storage systems in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port FICON Send I/O Rate (ops/s)**

The average number of send operations per second for Fibre Channel connection (FICON) usage.

**Port FICON Receive I/O Rate (ops/s)**

The average number of receive operations per second for FICON usage.

**Total Port FICON I/O Rate (ops/s)**

The average number of send and receive operations per second for FICON usage.

**Port FCP Send I/O Rate (ops/s)**

The average number of send operations per second for Fibre Channel Protocol (FCP) usage.

**Port FCP Receive I/O Rate (ops/s)**

The average number of receive operations per second for FCP usage.

**Total Port FCP I/O Rate (ops/s)**

The average number of send operations and receive operations per second for FCP usage.

**Port PPRC Send I/O Rate (ops/s)**

The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.

**Port PPRC Receive I/O Rate (ops/s)**

The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.

**Total Port PPRC I/O Rate (ops/s)**

The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port FICON Send Data Rate (MiB/s)**

The average number of MiB per second that is sent for FICON usage.

**Port FICON Receive Data Rate (MiB/s)**

The average number of MiB per second that is received for FICON usage.

**Total Port FICON Data Rate (MiB/s)**

The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.

**Port FCP Send Data Rate (MiB/s)**

The average number of MiB per second that are sent for FCP usage.

**Port FCP Receive Data Rate (MiB/s)**

The average number of MiB per second that are received for FCP usage.

**Total Port FCP Data Rate (MiB/s)**

The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

**Port PPRC Send Data Rate (MiB/s)**

The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.

**Port PPRC Receive Data Rate (MiB/s)**

The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.

**Total Port PPRC Data Rate (MiB/s)**

The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.

**Port Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation.

**Port Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation.

**Overall Port Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation.

**Port FCP Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation for FCP usage.

**Port FCP Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation for FCP usage.

**Overall Port FCP Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation for FCP usage.

**Port FICON Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation for FICON usage.

**Port FICON Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation for FICON usage.

**Overall Port FICON Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.

**Port PPRC Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.

**Port PPRC Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.

**Overall Port PPRC Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

**Port Send Transfer Size (KiB/op)**

The average number of KiB that are transferred per send operation.

**Port Receive Transfer Size (KiB/op)**

The average number of KiB that are transferred per receive operation.

**Overall Port Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

**Port Send Utilization Percentage**

The average percentage of time that the port is busy sending data.

**Port Receive Utilization Percentage**

The average percentage of time that the port is busy receiving data.

**Overall Port Utilization Percentage**

The average percentage of time that the port is busy sending or receiving data.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port error data**

You can create performance reports that include the following information:

**Error Frame Rate (count/s)**

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Out of Order Data Rate (count/s)**

The average number of times per second that an out-of-order frame is detected.

**Out of Order ACK Rate (count/s)**

The average number of times per second that an out-of-order acknowledge (ACK) frame is detected.

An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.

**Duplicate Frame Rate (count/s)**

The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.

**Invalid Relative Offset Rate (count/s)**

The average number of times per second that frames are received with an invalid relative offset in the frame header.

**Sequence Timeout Rate (count/s)**

The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.

**Link Reset Transmitted Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

**Link Reset Received Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

**Credit Recovery Link Reset Rate**

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

**Performance metrics for ports on an XIV system:**

You can include performance metrics for storage system ports on an XIV system in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.



**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation.

**Port Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation.

**Overall Port Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Performance metrics for storage system ports on block servers:**

You can include performance metrics for storage system ports on block servers in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation.

**Port Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation.

**Overall Port Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation.

**Port Send Transfer Size (KiB/op)**

The average number of KiB that are transferred per send operation.

**Port Receive Transfer Size (KiB/op)**

The average number of KiB that are transferred per receive operation.

**Overall Port Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

**Data for storage system RAID arrays in performance reports**

You can include general information, capacity data, and properties of storage system RAID arrays in performance reports.

**Information about storage system RAID arrays**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Controller Name**

The name that was assigned to the controller when it was added to the system.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage RAID Array Name**

The name that was assigned to the RAID array when it was added to the system.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage RAID Array DDM Capacity (GiB)**

The amount of storage space on the disk drive modules in an array.

**Component properties**

You can create performance reports that include the following information:

**Storage RAID Array RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage RAID Array Width**

The number of disks that are in the storage array.

**Storage RAID Array DDM Speed (RPM)**

The speed of the disk drive module in the array.

**Storage RAID Array Description**

Shows the array site that the array is associated with.

**Storage RAID Array Tag**

A number that identifies an array.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Performance metrics for storage system arrays:**

You can include front end data, cache data, back-end data, and other data for storage system arrays in performance reports.

**Front end data**

You can create performance reports that include the following information:

**Normal Read I/O Rate (ops/s)**

The average number of nonsequential read operations per second.

**Sequential Read I/O Rate (ops/s)**

The average number of sequential read operations per second.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Normal Write I/O Rate (ops/s)**

The average number of nonsequential write operations per second.

**Sequential Write I/O Rate (ops/s)**

The average number of sequential write operations per second.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Normal I/O Rate (ops/s)**

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

**Total Sequential I/O Rate (ops/s)**

The average number of sequential I/O operations per second. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Record Mode Read I/O Rate (ops/s)**

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Normal Read Cache Hit Percentage**

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Sequential Read Cache Hit Percentage**

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Normal Write Cache Hit Percentage**

The percentage of nonsequential write operations that are handled in the cache.

**Sequential Write Cache Hit Percentage**

The percentage of sequential write operations that are handled in the cache.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Normal Cache Hit Percentage**

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Sequential Cache Hit Percentage**

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Record Mode Read Cache Hit Percentage**

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Back-end data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Sequential Back-End I/O Percentage**

The percentage of all I/O operations that are issued to the back-end storage resources. The operations are for the array and are sequential operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Other data**

You can create performance reports that include the following information:

**HPF Read I/O Rate (ops/s)**

The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

**HPF Write I/O Rate (ops/s)**

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

**Total HPF I/O Rate (ops/s)**

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

**HPF I/O Percentage**

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

**PPRC Transfer Rate (ops/s)**

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

**Disk Utilization Percentage**

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed for this property. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

**Data for storage system controllers in performance reports**

You can include general information about storage system controllers in performance reports.

## Information about storage system controllers

You can create performance reports that include the following information:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

### Storage Controller Name

The name that was assigned to the controller when it was added to the system.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

## Performance metrics for storage system controllers on DS8000 and DS6000 systems:

You can include front end data, cache data, back-end data, and other data for storage system controllers on DS8000 and DS6000 series systems in performance reports.

### Front end data

You can create performance reports that include the following information:

#### Normal Read I/O Rate (ops/s)

The average number of nonsequential read operations per second.

#### Sequential Read I/O Rate (ops/s)

The average number of sequential read operations per second.

#### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

#### Normal Write I/O Rate (ops/s)

The average number of nonsequential write operations per second.

#### Sequential Write I/O Rate (ops/s)

The average number of sequential write operations per second.

#### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

#### Total Normal I/O Rate (ops/s)

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

#### Total Sequential I/O Rate (ops/s)

The average number of sequential I/O operations per second. This value includes both read and write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.



**Record Mode Read I/O Rate (ops/s)**

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Normal Read Cache Hit Percentage**

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Sequential Read Cache Hit Percentage**

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Normal Write Cache Hit Percentage**

The percentage of nonsequential write operations that are handled in the cache.

**Sequential Write Cache Hit Percentage**

The percentage of sequential write operations that are handled in the cache.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Normal Cache Hit Percentage**

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Sequential Cache Hit Percentage**

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Record Mode Read Cache Hit Percentage**

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Back-end data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Other data**

You can create performance reports that include the following information:

**HPF Read I/O Rate (ops/s)**

The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

**HPF Write I/O Rate (ops/s)**

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

**Total HPF I/O Rate (ops/s)**

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

**HPF I/O Percentage**

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

**PPRC Transfer Rate (ops/s)**

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

**Cache Holding Time (s)**

The average number of seconds that I/O data for a storage system controller is held in the cache. A short cache-holding time indicates adverse performance.

**Performance metrics for storage system controllers on block servers:**

You can include performance data for volumes and caches for storage system controllers on block servers in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Data for storage system modules in performance reports**

You can include properties of storage system modules, and other information about storage system modules in performance reports.

**Information about storage system modules**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Module Name**

The name that was assigned to the module when it was added to the system.

**Component properties**

You can create performance reports that include the following information:

**Storage Module IP Address**

The IP address of the resource.

**Storage Module WWN**

The worldwide name of the module.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Status information**

You can create performance reports that include the following information:

**Storage Module Status**

The condition of the resource, for example normal, warning, or error.

**Storage Module Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Performance metrics for storage system modules on an XIV system:**

You can include performance metrics for volumes, caches, and other data for storage system modules on an XIV system in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Read Data Cache Hit Percentage**

The percentage of all read data that is read from the cache.

**Write Data Cache Hit Percentage**

The percentage of all write data that is written to cache slots that are marked as modified.

**Overall Data Cache Hit Percentage**

The percentage of all data that is handled in the cache. This value includes

read data that is read from the cache and write data that is written to cache slots that are marked as modified.

**Read Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a read-cache hit operation.

**Write Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a write-cache hit operation.

**Overall Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

**Read Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a read-cache miss operation.

**Write Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a write-cache miss operation.

**Overall Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

**SSD Read Cache Hit Percentage**

The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**SSD Read Data Cache Hit Percentage**

The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**SSD Read Cache Hit Response Time (ms/op)**

The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

**Other data**

You can create performance reports that include the following information:

**Small Transfers I/O Percentage**

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.



**Medium Transfers I/O Percentage**

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers I/O Percentage**

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers I/O Percentage**

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

**Data for storage system pools in performance reports**

You can include general information, capacity data, properties, and other information about storage system pools in performance reports.

**Information about storage system pools**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage Pool Capacity (GiB)**

The total amount of storage space in a pool.

**Storage Pool Available Space (GiB)**

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$\text{pool capacity} - \text{used space}$

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

**Storage Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This

inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Pool Extent Size (MiB)**

The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Pool Number of Storage Volumes**

The number of volumes in the storage pool.

**Storage Pool Number of MDisks**

The number of managed disks in a storage pool.

**Storage Pool Physical Allocation Percentage**

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Pool Virtual Allocation Percentage**

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

### **Storage Pool Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{virtual capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB  $\div$  100 GiB  $\times$  100).

This value is only available for pools with thin-provisioned volumes.

### **Storage Pool Unallocated Volume Space (GiB)**

The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{total volume capacity} - \text{allocated space}$$

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

### **Storage Pool Unused Volume Space (GiB)**

The amount of space that is allocated to the volumes in a pool and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

### **Storage Pool Unallocatable Volume Space (GiB)**

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

**Storage Pool Used Space Percentage**

The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

**Storage Pool Compressed Virtual Capacity (GiB)**

The total virtual capacity of all the volumes that are compressed in a pool.

**Storage Pool Compressed Capacity (GiB)**

The amount of storage space that is used by compressed volumes in a pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

**Storage Pool Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

**Storage Pool Maximum I/O Capability**

The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Pool Back-End Storage System Type
- Storage Pool Back-End Storage RAID Level
- Storage Pool Back-End Storage Disk Type

This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Pool Back-End Storage Disks**

The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Component properties**

You can create performance reports that include the following information:

**Storage Pool Is Encryptable**

Shows whether the storage pool can be encrypted. If this value is Yes, the storage pool can be encrypted.

**Storage Pool Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

**Storage Pool Is Solid State**

Shows whether there are solid-state drives in the pool. This property can contain the following values:

**Mixed** The pool contains both hard disk drives and solid-state drives.

**Non solid state**

The pool contains no solid-state drives.

**Solid state**

The pool contains a least one solid-state drive.

**Storage Pool Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Pool RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Pool Warning Level**

The percentage of used capacity of the storage pool at which a warning is generated.

**Storage Pool Is Compression Active**

Shows whether the compression feature is enabled on the storage pool. If this value is Yes, the compression feature is enabled.

**Storage Pool Custom Tag 1, 2, and 3**

User-defined text that is associated with a storage pool. You can add or edit the custom tags for a storage pool on the **Properties** notebook of the pool in the Tivoli Storage Productivity Center web-based GUI.

**Storage Pool Back-End Storage System Type**

The type of storage system that provides storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Pool Back-End Storage RAID Level**

The RAID level of the volumes on the back-end storage system that provide storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Pool Back-End Storage Disk Type**

The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Status information**

You can create performance reports that include the following information:

**Storage Pool Status**

The condition of the resource, for example normal, warning, or error.

**Storage Pool Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

## **Performance metrics for pools on SAN Volume Controller and Storwize systems:**

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes and back-end arrays for pools on these systems in performance reports.

### **Volume data**

You can create performance reports that include the following information:

#### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

#### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

#### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### **Maximum Read I/O Rate (ops/s)**

The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

#### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

#### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

#### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

#### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

#### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

#### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### **Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

#### **Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

#### **Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

## Back-end array data

You can create performance reports that include the following information:

### **Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

### **Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

### **Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

### **Maximum Back-End Read I/O Rate (ops/s)**

The maximum number of read operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

### **Maximum Back-End Write I/O Rate (ops/s)**

The maximum number of write operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

### **Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

### **Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

### **Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

### **Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

### **Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

### **Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

### **Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

### **Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.



For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Maximum Back-End Read Response Time (ms/op)**

The maximum number of milliseconds for the back-end storage resources to respond to a read operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Maximum Back-End Write Response Time (ms/op)**

The maximum number of milliseconds for the back-end storage resources to respond to a write operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Performance metrics for pools on DS8000 and DS6000 storage systems:**

You can include performance metrics for volumes, caches, back-end arrays, and other data for pools that are on DS8000 and DS6000 storage systems in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Normal Read I/O Rate (ops/s)**

The average number of nonsequential read operations per second.

**Sequential Read I/O Rate (ops/s)**

The average number of sequential read operations per second.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Normal Write I/O Rate (ops/s)**

The average number of nonsequential write operations per second.

**Sequential Write I/O Rate (ops/s)**

The average number of sequential write operations per second.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Normal I/O Rate (ops/s)**

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

**Total Sequential I/O Rate (ops/s)**

The average number of sequential I/O operations per second. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Record Mode Read I/O Rate (ops/s)**

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Normal Read Cache Hit Percentage**

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Sequential Read Cache Hit Percentage**

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Normal Write Cache Hit Percentage**

The percentage of nonsequential write operations that are handled in the cache.

**Sequential Write Cache Hit Percentage**

The percentage of sequential write operations that are handled in the cache.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Normal Cache Hit Percentage**

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Sequential Cache Hit Percentage**

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Record Mode Read Cache Hit Percentage**

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

**Other data**

You can create performance reports that include the following information:

**HPF Read I/O Rate (ops/s)**

The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

**HPF Write I/O Rate (ops/s)**

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

**Total HPF I/O Rate (ops/s)**

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

**HPF I/O Percentage**

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

**PPRC Transfer Rate (ops/s)**

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

## Performance metrics for pools on an XIV system:

You can include performance data for volumes, caches, and other data for pools on an XIV system in performance reports.

### Volume data

You can create performance reports that include the following information:

#### Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

#### Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

#### Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

#### Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations.

#### Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations.

#### Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

#### Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

#### Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

#### Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

#### Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

#### Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

#### Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

### Cache data

You can create performance reports that include the following information:

#### Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Read Data Cache Hit Percentage**

The percentage of all read data that is read from the cache.

**Write Data Cache Hit Percentage**

The percentage of all write data that is written to cache slots that are marked as modified.

**Overall Data Cache Hit Percentage**

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

**Read Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a read-cache hit operation.

**Write Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a write-cache hit operation.

**Overall Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

**Read Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a read-cache miss operation.

**Write Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a write-cache miss operation.

**Overall Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

**Other data**

You can create performance reports that include the following information:

**Small Transfers I/O Percentage**

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers I/O Percentage**

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers I/O Percentage**

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers I/O Percentage**

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

**Data for managed disks on storage systems in performance reports**

You can include general information, capacity data, properties, and other information about managed disks on storage systems in performance reports.

**Information about managed disks on storage systems**

You can create performance reports that include the following information:



**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Storage System MDisk Name**

The name that was assigned to the managed disk on a storage system when it was added to the system.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage MDisk Available Space (GiB)**

The amount of unused storage space on the managed disk on a storage system.

**Storage System MDisk Strip Size (KB)**

The RAID strip size on a managed disk on a storage system.

**Component properties**

You can create performance reports that include the following information:

**Storage MDisk RAID Level**

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

**Storage MDisk Spare Goal**

The number of spare drives that are required to maintain redundancy. Use spare drives to protect the system against drive failures in the array on a managed disk on a storage virtualizer.

**Storage System MDisk Type**

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

**Storage System MDisk Mode**

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

**Array** Extents are provided from local disks.

**Managed**

Extents are provided from a back-end storage volume.

**Unmanaged**

The managed disk is not used in the system.

**Storage System MDisk Is Balanced**

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is Yes, the LUNs are balanced.

**Storage System MDisk Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

**Storage System MDisk Write Verify**

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is Yes, all write operations are verified by a follow-up read operation.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Status information**

You can create performance reports that include the following information:

**Storage MDisk Status**

The condition of the resource, for example normal, warning, or error.

**Storage MDisk Native Status**

Shows the level of access of nodes in the system to a managed disk or volume on a storage system. The level of access can be online, offline, degraded, or excluded.

**Performance metrics for managed disks on SAN Volume Controller and Storwize systems:**

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for back-end arrays on managed disks on these systems in performance reports.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Data for storage system nodes in performance reports**

You can include general information, capacity data, properties, and other information about storage system nodes in performance reports.

**Information about storage system nodes**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Node Name**

The name that was assigned to the storage system node when it was added to the system.

**Storage I/O Group Name**

The name that was assigned to the I/O group when it was added to the system.

**Component properties**

You can create performance reports that include the following information:

**Storage Node IP Address**

The IP address of the resource.

**Storage Node WWN**

The worldwide name of the node.

**Storage Node Configuration**

The number of the node for the resource. This value can be either 0 or 1.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Status information**

You can create performance reports that include the following information:

**Storage Node Status**

Shows whether a storage system node is online.

**Performance metrics for nodes on SAN Volume Controller and Storwize systems:**

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include

performance metrics for ports, front ends, caches, back-end arrays, and other data for nodes on these systems in performance reports.

### **Port data**

You can create performance reports that include the following information:

#### **Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

#### **Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

#### **Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

#### **Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

#### **Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

#### **Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

#### **Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

#### **Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

#### **Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

#### **Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

#### **Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

#### **Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

#### **Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

#### **Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Local Node Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Remote Node Response Time (ms/op)**

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

**Overall Port to Local Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the

queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

**Overall Port to Remote Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

**Front end data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.



## Cache data

You can create performance reports that include the following information:

### **Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

### **Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

### **Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

### **Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

### **Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### **Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

### **Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

### **Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

### **Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

### **Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

### **Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Back-end data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

### **Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

### **Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

### **Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

### **Zero Buffer-to-Buffer Credit Percentage**

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

## **Data for storage system I/O groups in performance reports**

You can include general information, capacity data, properties, and other information about storage system I/O groups in performance reports.

### **Information about storage system I/O groups**

You can create performance reports that include the following information:

#### **Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

#### **Storage I/O Group Name**

The name that was assigned to the I/O group when it was added to the system.

### **Component properties**

You can create performance reports that include the following information:

#### **Storage I/O Group Mirroring Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide volume mirroring functions.

#### **Storage I/O Group Mirroring Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide volume mirroring functions.

**Storage I/O Group FlashCopy Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide FlashCopy mirroring functions.

**Storage I/O Group FlashCopy Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide FlashCopy mirroring functions.

**Storage I/O Group Remote Copy Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide mirroring functions by using a remote copy.

**Storage I/O Group Remote Copy Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide mirroring functions by using a remote copy.

**Storage I/O Group RAID Total Memory (MiB)**

The total amount of memory that is allocated on the nodes in an I/O group to provide RAID mirroring functions.

**Storage I/O Group RAID Free Memory (MiB)**

The amount of available memory on the nodes in an I/O group to provide RAID mirroring functions.

**Status information**

You can create performance reports that include the following information:

**Storage I/O Group Maintenance**

Shows whether the I/O group is in maintenance mode. If this value is Yes, the I/O group is in maintenance mode.

**Storage I/O Group Is Compression Active**

Shows whether the compression feature is enabled on the volumes in the storage I/O group. If this value is Yes, the compression feature is enabled.

**Storage I/O Group Is Compression Supported**

Shows whether the compression feature is available for volumes in the storage I/O group. If this value is Yes, the compression feature is available.

**Storage System Configuration**

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

**Performance metrics for I/O groups on SAN Volume Controller and Storwize systems:**

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, volumes, caches, back-end arrays, and other data for I/O groups on these systems in performance reports.

**Port data**

You can create performance reports that include the following information:

**Port Send I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which data is sent from a port.

**Port Receive I/O Rate (ops/s)**

The average number of I/O operations per second for operations in which the port receives data.

**Total Port I/O Rate (ops/s)**

The average number of send operations and receive operations per second.

**Port to Host Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to host computers.

**Port to Host Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from host computers.

**Total Port to Host I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between host computers and the component.

**Port to Disk Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to back-end storage resources.

**Port to Disk Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from back-end storage resources.

**Total Port to Disk I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

**Port to Local Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

**Total Port to Local Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send I/O Rate (ops/s)**

The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive I/O Rate (ops/s)**

The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

**Total Port to Remote Node I/O Rate (ops/s)**

The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Host Send Data Rate (MiB/s)**

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

**Port to Host Receive Data Rate (MiB/s)**

The average rate at which data is received from host computers. The rate is measured in MiB per second.

**Total Port to Host Data Rate (MiB/s)**

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Disk Send Data Rate (MiB/s)**

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

**Port to Disk Receive Data Rate (MiB/s)**

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

**Total Port to Disk Data Rate (MiB/s)**

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

**Port to Local Node Send Data Rate (MiB/s)**

The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Receive Data Rate (MiB/s)**

The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Total Port to Local Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Send Data Rate (MiB/s)**

The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Remote Node Receive Data Rate (MiB/s)**

The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.



**Total Port to Remote Node Data Rate (MiB/s)**

The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

**Port to Local Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Local Node Response Time (ms/op)**

The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Send Response Time (ms/op)**

The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Remote Node Receive Response Time (ms/op)**

The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Overall Port to Remote Node Response Time (ms/op)**

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

**Port to Local Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

**Port to Local Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

**Overall Port to Local Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

**Port to Remote Node Send Queue Time (ms/op)**

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

**Port to Remote Node Receive Queue Time (ms/op)**

The average time in milliseconds that a receive operation spends in the

queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

**Overall Port to Remote Node Queue Time (ms/op)**

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Back-end array data**

You can create performance reports that include the following information:

**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (op/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Processor Utilization Percentage**

The average percentage of time that the processors on nodes are busy.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume

Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

**Nonpreferred Node Usage Percentage**

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Zero Buffer-to-Buffer Credit Timer (microseconds)**

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

**Zero Buffer-to-Buffer Credit Percentage**

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

**Data for local disks on storage systems in performance reports**

You can include general information, capacity data, properties, and other information about local disks that are on storage systems in performance reports.

**Information about local disks on storage systems**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage MDisk Name**

The name that was assigned to the managed disk when it was added to the system.

**Storage Local Disk Name**

The name that was assigned to the local disk on the storage system when the local disk was created on the system.

**Capacity and usage data**

You can create performance reports that include the following information:

**Storage Local Disk Capacity (GiB)**

The amount of storage space on the local disk.

## Component properties

You can create performance reports that include the following information:

### Storage Local Disk Is Encryptable

Shows whether the resource can be encrypted. If this value is Yes, the resource can be encrypted.

### Storage Local Disk Is Encrypted

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

### Storage Local Disk Storage Class

The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

### Storage Local Disk Is Solid State

Shows whether the disk is a solid-state drive.

### Storage Local Disk Speed (RPM)

The speed of the local disk on a storage system.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is Storage System, the storage system is configured as back-end storage. If this value is Storage Virtualizer, the storage system is configured to virtualize back-end storage.

## Status information

You can create performance reports that include the following information:

### Storage Local Disk Status

The condition of the resource, for example normal, warning, or error.

## Vendor, model, and device information

You can create performance reports that include the following information:

### Storage System Local Disk Firmware Version

The version number of the firmware that is running on the disk.

### Storage System Local Disk Model

The model name or model number of the resource.

### Storage System Local Disk Serial Number

The serial number of the resource.

### Storage System Local Disk Vendor

The vendor who supplied the resource.

## Performance metrics for local disks on storage systems:

You can include performance data for back-end arrays on local disks that are on storage systems in performance reports.

### Back-end array data

You can create performance reports that include the following information:



**Back-End Read I/O Rate (ops/s)**

The average number of read operations per second that are issued to the back-end storage resources.

**Back-End Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the back-end storage resources.

**Total Back-End I/O Rate (ops/s)**

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

**Back-End Read Data Rate (MiB/s)**

The average number of MiB per second that are read from the back-end storage resources.

**Back-End Write Data Rate (MiB/s)**

The average number of MiB per second that are written to the back-end storage resources.

**Total Back-End Data Rate (MiB/s)**

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

**Back-End Read Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Back-End Write Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Overall Back-End Response Time (ms/op)**

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

**Peak Back-End Read Response Time (ms)**

The longest time for a back-end storage resource to respond to a read operation.

**Peak Back-End Write Response Time (ms)**

The longest time for a back-end storage resource to respond to a write operation by a node.

**Back-End Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation from the back-end storage resources.

**Back-End Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation to the back-end storage resources.

**Overall Back-End Transfer Size (KiB/op)**

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

**Back-End Read Queue Time (ms/op)**

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Back-End Write Queue Time (ms/op)**

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Overall Back-End Queue Time (ms/op)**

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Read Queue Time (ms)**

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

**Peak Back-End Write Queue Time (ms)**

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

**Data for storage host connections in performance reports**

You can include general information, capacity data, properties, and other information about storage host connections in performance reports.

**Information about storage host connections**

You can create performance reports that include the following information:

**Storage System Name**

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

**Storage Host Connection Name**

The user-defined name that describes the server or storage device that is assigned to a volume.

**XIV Cluster**

The name of the cluster, if any, as defined on the XIV system. If no cluster is defined for the host connection, this field is blank.

**Performance metrics for host connections on SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified systems:**

You can include performance metrics for volumes, caches, and for other data for host connections in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Write Cache Flush Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in flush-through mode.

**Write Cache Write Through I/O Rate (ops/s)**

The average number of tracks per second that are written to disk in write-through mode.

**Write Cache Overflow I/O Rate (ops/s)**

The average number of write operations per second that are delayed because there is not enough space in the write cache.

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Read Ahead Percentage of Cache Hits**

The percentage of all read-cache hits that occur on pre-staged data.

**Dirty Write Percentage of Cache Hits**

The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

**Write Cache Flush Through Percentage**

The percentage of tracks that is written to disk in flush-through mode.

**Write Cache Write Through Percentage**

The percentage of tracks that is written to disk in write-through mode.

**Write Cache Overflow Percentage**

The percentage of write operations that are delayed because there is not enough space in the write cache.

**Other data**

You can create performance reports that include the following information:

**Global Mirror Write I/O Rate (ops/s)**

The average number of write operations per second that are issued to the Global Mirror secondary site.

**Global Mirror Overlapping Write I/O Rate (op/s)**

The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Peak Read Response Time (ms)**

The longest time to complete a read operation.

**Peak Write Response Time (ms)**

The longest time to complete a write operation.

**Global Mirror Secondary Write Lag (ms/op)**

The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

**Overall Host Attributed Response Time Percentage**

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

**Global Mirror Overlapping Write Percentage**

The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

**Nonpreferred Node Usage Percentage**

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

**Performance metrics for host connections on DS8000 and DS6000 storage systems:**

You can include performance metrics for volumes, caches, and for other data for storage host connections on DS8000 and DS6000 storage systems in performance reports.

**Volume data**

You can create performance reports that include the following information:

**Normal Read I/O Rate (ops/s)**

The average number of nonsequential read operations per second.

**Sequential Read I/O Rate (ops/s)**

The average number of sequential read operations per second.

**Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

**Normal Write I/O Rate (ops/s)**

The average number of nonsequential write operations per second.

**Sequential Write I/O Rate (ops/s)**

The average number of sequential write operations per second.

**Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

**Total Normal I/O Rate (ops/s)**

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

**Total Sequential I/O Rate (ops/s)**

The average number of sequential I/O operations per second. This value includes both read and write operations.

**Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

**Record Mode Read I/O Rate (ops/s)**

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

**Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

**Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

**Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

**Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

**Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

**Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

**Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

**Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

**Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

**Cache data**

You can create performance reports that include the following information:

**Write Cache Delay I/O Rate (ops/s)**

The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Normal Read Cache Hit Percentage**

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Sequential Read Cache Hit Percentage**

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Normal Write Cache Hit Percentage**

The percentage of nonsequential write operations that are handled in the cache.

**Sequential Write Cache Hit Percentage**

The percentage of sequential write operations that are handled in the cache.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Normal Cache Hit Percentage**

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Sequential Cache Hit Percentage**

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Record Mode Read Cache Hit Percentage**

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

**Disk to Cache Transfer Rate (ops/s)**

The average number of tracks that are transferred per second from the disks to the cache.

**Cache to Disk Transfer Rate (ops/s)**

The average number of tracks per second that are transferred from the cache to the disks.

**Write Cache Delay Percentage**

The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

**Other data**

You can create performance reports that include the following information:

**HPF Read I/O Rate (ops/s)**

The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

**HPF Write I/O Rate (ops/s)**

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

**Total HPF I/O Rate (ops/s)**

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

**HPF I/O Percentage**

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

**PPRC Transfer Rate (ops/s)**

The average number of tracks per second that are transferred to the



secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

### **Performance metrics for storage host connections on an XIV system:**

You can include performance metrics for volumes, caches, and for other data for storage host connections on an XIV system in performance reports.

#### **Volume data**

You can create performance reports that include the following information:

##### **Overall Read I/O Rate (ops/s)**

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

##### **Overall Write I/O Rate (ops/s)**

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

##### **Total Overall I/O Rate (ops/s)**

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

##### **Read Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations.

##### **Write Data Rate (MiB/s)**

The average number of MiB per second that are transferred for write operations.

##### **Total Data Rate (MiB/s)**

The average number of MiB per second that are transferred for read operations and write operations.

##### **Read Response Time (ms/op)**

The average number of milliseconds to complete a read operation.

##### **Write Response Time (ms/op)**

The average number of milliseconds to complete a write operation.

##### **Overall Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

##### **Read Transfer Size (KiB/op)**

The average number of KiB that are transferred per read operation.

##### **Write Transfer Size (KiB/op)**

The average number of KiB that are transferred per write operation.

##### **Overall Transfer Size (KiB/op)**

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

#### **Cache data**

You can create performance reports that include the following information:

**Overall Read Cache Hit Percentage**

The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Overall Write Cache Hit Percentage**

The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

**Total Overall Cache Hit Percentage**

The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

**Read Data Cache Hit Percentage**

The percentage of all read data that is read from the cache.

**Write Data Cache Hit Percentage**

The percentage of all write data that is written to cache slots that are marked as modified.

**Overall Data Cache Hit Percentage**

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

**Read Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a read-cache hit operation.

**Write Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a write-cache hit operation.

**Overall Cache Hit Response Time (ms/op)**

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

**Read Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a read-cache miss operation.

**Write Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a write-cache miss operation.

**Overall Cache Miss Response Time (ms/op)**

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

**Other data**

You can create performance reports that include the following information:

**Small Transfers I/O Percentage**

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers I/O Percentage**

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers I/O Percentage**

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers I/O Percentage**

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Data Percentage**

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

**Small Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

**Medium Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

**Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

**Very Large Transfers Response Time (ms/op)**

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

**Data for switches in performance reports**

You can include properties and other information about switches in performance reports.

## Information about switches

You can create performance reports that include the following information:

### Switch Logical Name

The logical name of a switch or switch port.

## Component properties

You can create performance reports that include the following information:

### Switch WWN

The World Wide Name (WWN) of the switch. A WWN is the unique 64-bit identifier for the switch.

### Parent Switch WWN

The worldwide name of the parent switch.

### Switch Management Telnet Address

The Telnet address that you access to manage the switch.

### Switch Management SNMP Address

The Simple Network Management Protocol (SNMP) address that you access to manage the switch.

### Switch Management URL Address

The URL that you access to manage the switch.

### Switch Domain

The domain ID of a switch. The ID is an 8-bit identifier with a range of 0-255. This column is blank for physical switches that are parents of virtual switches.

### Switch Version

The version number of the operating system that is running on the switch.

### Switch Location

The physical location of the switch. The location is defined when a switch is added to Tivoli Storage Productivity Center. In the web-based GUI, you can add or edit the location of the switch on the **General** tab of the properties notebook.

### Switch Custom Tag 1, 2, and 3

User-defined text that is associated with the switch. In the web-based GUI, you can add or edit the custom tags for the switch on the **General** tab of the properties notebook.

## Status information

You can create performance reports that include the following information:

### Switch Status

The condition of the resource, for example normal, warning, or error.

### Switch Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

## Vendor, model, and device information

You can create performance reports that include the following information:

**Switch IP Address**

The IP address of the resource.

**Switch Serial Number**

The serial number of the resource.

**Switch Vendor**

The vendor who supplied the resource.

**Switch Model**

The model name or model number of the resource.

**Performance metrics for switches:**

You can include performance data for switches, including switch error data, in performance reports.

**Switch data**

You can create performance reports that include the following information:

**Port Send Frame Rate (frames/s)**

The average number of frames per second that are sent through the port.

**Port Receive Frame Rate (frames/s)**

The average number of frames per second that are received by the port.

**Total Port Frame Rate (frames/s)**

The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port Send Frame Size (KiB/op)**

The average size of a frame, in KiB, that is sent through the port.

**Port Receive Frame Size (KiB/op)**

The average size of a frame, in KiB, that is received by the port.

**Overall Port Frame Size (KiB/op)**

The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Switch error data**

You can create performance reports that include the following information:

**Error Frame Rate (count/s)**

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

**Discarded Frame Rate (count/s)**

The average number of frames per second that are discarded because host buffers are unavailable for the port.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Encoding Disparity Error Rate (count/s)**

The average number of disparity errors per second that are received.

**Discarded Class 3 Frame Rate (count/s)**

The average number of class 3 frames per second that are discarded.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Link Reset Transmitted Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

**Link Reset Received Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

**Short Frame Rate (count/s)**

The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

**Long Frame Rate (count/s)**

The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.

**F-BSY Frame Rate (count/s)**

The average number of F-BSY frames per second that are generated.

An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N\_port is busy.

**F-RJT Frame Rate (count/s)**

The average number of F-RJT frames per second that are generated.

An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

**Zero Buffer Credit Rate**

The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

**Class 3 Send Timeout Frame Rate**

The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

**Class 3 Receive Timeout Frame Rate**

The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

#### **RDY Priority Override Rate**

The average number of times per second during which the sending of R\_RDY or VC\_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

#### **Port State Change Rate**

The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

#### **Port Congestion Index**

The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

#### **Zero Buffer-to-Buffer Credit Percentage**

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

#### **Credit Recovery Link Reset Rate**

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

### **Data for switch ports in performance reports**

You can include properties and other information about switch ports in performance reports.

#### **Information about switch ports**

You can create performance reports that include the following information:

##### **Switch Logical Name**

The logical name of a switch or switch port.

##### **Switch Port WWPN**

The worldwide port name of the port on the switch.

#### **Component properties**

You can create performance reports that include the following information:



**Switch Port Number**

The port number on the resource.

**Switch Port Type**

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N\_Port, F\_Port, or another type of port.

**Switch Port Speed (GiB/s)**

The speed of a port, which is measured in GiB per second.

**Switch Blade Slot Number**

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

**Status information**

You can create performance reports that include the following information:

**Switch Port Status**

The condition of the resource, for example normal, warning, or error.

**Switch Port Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Switch Port Enabled State**

Shows whether a port is enabled, disabled, or is enabled but offline.

**Switch Port Operational Status**

The current operational state of the switch port. For example, this value can be Error or OK.

**Performance metrics for switch ports:**

You can include performance data for switch ports, including port error data, in performance reports.

**Switch port data**

You can create performance reports that include the following information:

**Port Send Frame Rate (frames/s)**

The average number of frames per second that are sent through the port.

**Port Receive Frame Rate (frames/s)**

The average number of frames per second that are received by the port.

**Total Port Frame Rate (frames/s)**

The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

**Port Send Data Rate (MiB/s)**

The average rate at which data is sent through the port. The rate is measured in MiB per second.

**Port Receive Data Rate (MiB/s)**

The average rate at which data is received by the port. The rate is measured in MiB per second.

**Total Port Data Rate (MiB/s)**

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

**Port Peak Send Data Rate (MiB/s)**

The fastest rate at which data is sent through this switch port.

**Port Peak Receive Data Rate (MiB/s)**

The fastest rate at which data is received through this switch port.

**Port Send Frame Size (KiB/op)**

The average size of a frame, in KiB, that is sent through the port.

**Port Receive Frame Size (KiB/op)**

The average size of a frame, in KiB, that is received by the port.

**Overall Port Frame Size (KiB/op)**

The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.

**Port Send Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Port Receive Bandwidth Percentage**

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Overall Port Bandwidth Percentage**

The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

**Data for errors on switch ports**

You can create performance reports that include the following information:

**Error Frame Rate (count/s)**

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

**Discarded Frame Rate (count/s)**

The average number of frames per second that are discarded because host buffers are unavailable for the port.

**Link Failure Rate (count/s)**

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

**Loss of Sync Rate (count/s)**

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

**Loss of Signal Rate (count/s)**

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused

by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

**CRC Error Rate (count/s)**

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

**Encoding Disparity Error Rate (count/s)**

The average number of disparity errors per second that are received.

**Discarded Class 3 Frame Rate (count/s)**

The average number of class 3 frames per second that are discarded.

**Primitive Sequence Protocol Error Rate (count/s)**

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

**Invalid Transmission Word Rate (count/s)**

The average number of bit errors per second that are detected.

**Link Reset Transmitted Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

**Link Reset Received Rate (count/s)**

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

**Short Frame Rate (count/s)**

The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

**Long Frame Rate (count/s)**

The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.

**F-BSY Frame Rate (count/s)**

The average number of F-BSY frames per second that are generated.

An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N\_port is busy.

**F-RJT Frame Rate (count/s)**

The average number of F-RJT frames per second that are generated.

An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

**Zero Buffer Credit Rate**

The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to

send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

### **Class 3 Send Timeout Frame Rate**

The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

### **Class 3 Receive Timeout Frame Rate**

The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

### **RDY Priority Override Rate**

The average number of times per second during which the sending of R\_RDY or VC\_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

### **Port State Change Rate**

The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

### **Port Congestion Index**

The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

### **Zero Buffer-to-Buffer Credit Percentage**

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

### Credit Recovery Link Reset Rate

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

---

## Custom reports about capacity and relationships

You can use Tivoli Common Reporting to create custom reports about capacity and relationships. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources.

Capacity and relationship reports can contain information about the following resources and their components:

- Fabrics
- Switches
- Servers
- Hypervisors
- Network-attached storage servers
- Storage virtualizers
- Storage systems
- Storage resource groups
- Groups

### Creating custom capacity and relationship reports

You can create custom reports that show capacity and relationship information for resources. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources. You access custom capacity and relationship reports from the Tivoli Storage Productivity Center web-based GUI, and you create the reports in the IBM Cognos reporting software.


Before you can create a custom report about capacity and relationships, ensure that data was collected from the resources. The following table shows where the data must be collected from for each resource:

Resource	Prerequisite
Fabric	Data must be collected from the fabric by a CIM agent, SNMP agent, or Storage Resource agent.
Switch	Data must be collected from the switch by a CIM agent, SNMP agent, or Storage Resource agent.
Server	Data must be collected from the server by the Storage Resource agent, or by an agentless server.
Hypervisor	Data must be collected from the VMware ESX server.

Resource	Prerequisite
NAS	Data must be collected from any NetApp or IBM Scale Out Network Attached Storage resources that are in your storage environment.
Storage virtualizer	Data must be collected from the any of the following storage virtualizers that are in your storage environment: <ul style="list-style-type: none"> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> <li>• SAN Volume Controller</li> <li>• Hitachi Universal Storage Platform V</li> </ul>
Storage system	Data must be collected from the storage system, or any of the following systems that are configured as back-end storage: <ul style="list-style-type: none"> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> <li>• SAN Volume Controller</li> <li>• Hitachi Universal Storage Platform V</li> </ul>

**Restriction:** If you create a custom report that includes properties from both **Servers to Storage Systems** and **Fabrics and Switches**, the relationships between resources are not displayed correctly.

**Tip:** Before you run a report, you can preview the report with no data.


1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the type of report that you want to see, complete one of the following steps:
  - To create a basic report with limited formatting about the capacity and relationships of resources, select **Cognos > Create capacity reports**.
  - To create a report about the performance of resources and apply advanced formatting features to the report, complete the following steps:
    - a. Select **Cognos > Create capacity reports**.
    - b. Click the **Return** icon .
    - c. Click **Launch > Report Studio**.
3. Expand **Tivoli Storage Productivity Center** in the list.
4. Expand **Servers to Storage Systems** or **Fabrics and Switches** in the list.
5. Explore the folders in the list to locate the resources, properties, and statistics that you want to include in the report.
6. Depending on your browser, do one of the following steps to add properties and statistics to the report:
  - In Internet Explorer, drag the items that you want to see information about to the work area.
  - In Firefox, select the items that you want to see information about, and then click **Insert**.
7. Click **Save**.

**Related tasks:**

“Adding filters to reports”

You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.

**Related information:**

 [Run a Report](#)

For more information about how to preview and run reports, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cr\\_rptstd.10.2.0.doc/t\\_cr\\_rptstd\\_intrd\\_run\\_report.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cr_rptstd.10.2.0.doc/t_cr_rptstd_intrd_run_report.html).

## Adding filters to reports

You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.

Filters for resources are in the same list as the resources, in the appropriate resource folder. For example, to view the filters for servers, select **Server Resources > Server Resource Filters > Server Filters**.

Filters are available for all resources, except network-attached storage resources.

1. To apply a filter to the data in the work area, locate the filter that you want to include in the folders in the list.
2. Depending on your browser, do one of the following steps to add a filter to the report:
  - In Internet Explorer, drag the filter that you want to use to the work area.
  - In Firefox, select the filter that you want to use, and then click **Insert**.

The filter is applied to the report. If you did not add properties and statistics to the report, the filter is applied when you add the properties and statistics.

3. Optional: To delete the filter, right-click the filter in the work area, and then select **Delete**.

**Related information:**

 [Create a Simple Filter](#)

For more information about how to use filters with reports, go to the product documentation at [http://www.ibm.com/support/knowledgecenter/SSEP7J\\_10.2.0/com.ibm.swg.ba.cognos.ug\\_cr\\_qstd.10.2.0.doc/t\\_create\\_a\\_simple\\_filter.html](http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.0/com.ibm.swg.ba.cognos.ug_cr_qstd.10.2.0.doc/t_create_a_simple_filter.html).

### Filters for servers in capacity and relationship reports

You can filter capacity and relationship reports to include only the servers, server file systems, and server logical volumes that meet specific criteria.

#### Filters for servers

You can create capacity and relationship reports that include the following filters:

##### **Servers Connected to Virtualizers or Storage Systems**

Shows servers that are connected to storage virtualizers or storage systems. You can use this filter to identify servers that are connected directly to storage virtualizers, or connected directly or indirectly to storage systems. This filter excludes local disks from the report.

##### **Servers Directly Connected to Storage Systems**

Shows servers that are directly connected to a back-end storage system.

You can use this filter to identify servers or server disks that are connected to a storage system, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or server disks that are not virtualized. These file systems or server disks might be affected if a storage system is removed.

#### **Servers Directly Connected to Virtualizers**

Shows servers that are directly connected to a storage virtualizer. You can use this filter to identify servers that are connected to a storage virtualizer, but are not connected directly to a back-end storage system.

For example, you can use this filter to identify only server disks that are on a SAN Volume Controller.

### **Filters for server file systems**

You can create capacity and relationship reports that include the following filters:

#### **Rows with File System Mount Points**

Shows the file systems that have a mount point on the server.

#### **Rows without File System Mount Points**

Shows resources that do not have a mount point on the server. You can use this filter to identify disks that do not have file systems, that is, disks that are not in use.

For example, to identify unused space include the Server Name property and the Server Disk Path property in a report, and then include the Rows without File System Mount Points filter.

### **Filters for server logical volumes**

You can create capacity and relationship reports that include the following filters:

#### **Rows with Logical Volume Path**

Shows the logical volumes that have a path to the server.

#### **Rows without Logical Volume Path**

Shows resources that do not have a logical volume path to the server. You can use this filter to identify volumes that do not have paths, that is, volumes that are not in use.

For example, to identify unused space include the Server Name property and the Server Disk Path property in a report, and then include the Rows without Logical Volume Path filter.

### **Filters for hypervisors in capacity and relationship reports**

You can filter capacity and relationship reports to include only the hypervisors, hypervisor file systems, and hypervisor logical volumes that meet specific criteria.

### **Filters for hypervisors**

You can create capacity and relationship reports that include the following filters:

#### **Hypervisor with Probed Servers**

Shows hypervisors that are connected to servers. You can use this filter to include only hypervisors and their associated servers from which data was collected. For example, you might collect data only from your critical



production servers on a hypervisor, but not your test servers. In this case, you can use this filter to exclude information for the test servers from the report.

#### **Hypervisors Connected to Virtualizers or Storage Systems**

Shows hypervisors that are connected to storage virtualizers or storage systems. You can use this filter to identify hypervisors that are connected directly to storage virtualizers, or connected directly or indirectly to storage systems. This filter excludes local disks from the report.

#### **Hypervisors Directly Connected to Storage Systems**

Shows hypervisors that are directly connected to a back-end storage system. You can use this filter to identify hypervisors or hypervisor disks that are connected to a storage system, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or hypervisor disks that are not virtualized. These file systems or hypervisor disks might be affected if a storage system is removed.

#### **Hypervisors Directly Connected to Virtualizers**

Shows hypervisors that are directly connected to a storage virtualizer. You can use this filter to identify hypervisors that are connected to a storage virtualizer, but are not connected directly to a back-end storage system.

For example, you can use this filter to identify only hypervisor disks that are on a SAN Volume Controller.

### **Filters for hypervisor file systems**

You can create capacity and relationship reports that include the following filters:

#### **Hypervisors with File System Mount Points**

Shows the file systems that have a mount point on the hypervisor.

#### **Hypervisors without File System Mount Points**

Shows resources that do not have a mount point on the hypervisor. You can use this filter to identify disks that do not have file systems, that is, disks that are not in use.

For example, to identify unused space include the Hypervisor Name property and the Hypervisor Disk Path property in a report, and then include the Hypervisors without File System Mount Points filter.

### **Filters for hypervisor logical volumes**

You can create capacity and relationship reports that include the following filters:

#### **Hypervisors with Logical Volume Path**

Shows the logical volumes that have a path to the hypervisor.

#### **Hypervisors without Logical Volume Path**

Shows resources that do not have a logical volume path to the hypervisor. You can use this filter to identify volumes that do not have paths, that is, volumes that are not in use.

For example, to identify unused space include the Hypervisor Name property and the Hypervisor Disk Path property in a report, and then include the Hypervisors without Logical Volume Path filter.

## **Filters for storage virtualizers in capacity and relationship reports**

You can filter capacity and relationship reports to include only the storage virtualizers and storage virtualizer volumes that meet specific criteria.

### **Storage virtualizer filters**

You can create capacity and relationship reports that include the following filters:

#### **Virtualizers Not Connected to Servers**

Shows storage virtualizers that are not connected to servers.

### **Filters for storage virtualizer volumes**

You can create capacity and relationship reports that include the following filters:

#### **Rows with Virtualizer Volumes**

Shows storage virtualizer disks that are connected to the storage virtualizer that have an associated volume.

#### **Rows without Virtualizer Volumes**

Shows storage virtualizer disks that are connected to the storage virtualizer that do not have an associated volume.

#### **Virtualizer Primary Volume Information**

Shows primary volumes that are connected to the storage virtualizer. You can use this filter if you want to get a total of the used space on primary volumes in your storage environment.

#### **Virtualizer Volume Copy Information**

Shows copy volumes that are connected to the storage virtualizer. You can use this filter if you want to get a total of the used space on volume copies in your storage environment.

## **Filters for storage systems in capacity and relationship reports**

You can filter capacity and relationship reports to include only the storage systems or storage volumes that meet specific criteria.

### **Filters for storage systems**

You can create capacity and relationship reports that include the following filters:

#### **Storage Systems Directly Connected to Servers**

Shows back-end storage systems that are directly connected to a server. You can use this filter to identify storage systems that are connected to a server, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or server disks that are not virtualized. These file systems or server disks might be affected if a storage system is removed.

#### **Storage Systems Connected to Servers through Virtualizers**

Shows storage systems that are indirectly connected to a server through a storage virtualizer. You can use this filter to identify storage systems that are connected to a storage virtualizer, but are not connected directly to a server.

For example, you can use this filter to identify only server disks that are on a SAN Volume Controller.

**Storage Systems not Connected to Servers**

Shows storage systems that are not connected to a server either directly, or indirectly through a storage virtualizer.

**Rows with Storage Volumes**

Shows storage resources that have an associated storage volume.

**Rows without Storage Volumes**

Shows storage resources that do not have an associated storage volume.

**Filters for storage resource groups in capacity and relationship reports**

You can filter capacity and relationship reports to include only the storage resource groups that meet specific criteria.

**Filters for storage resource groups**

You can create capacity and relationship reports that include the following filters:

**Storage Resource Group for Volumes**

Shows only the volumes that are in a storage resource group.

**Storage Resource Group for Pools**

Shows only the pools that are in a storage resource group.

**Data and properties in capacity and relationship reports**

You can include capacity data, properties, relationship information, and other information about resources in your storage environment in capacity and relationship reports.

**Data for clusters in capacity and relationship reports**

You can include general information about clusters in capacity and relationship reports.

**Information about clusters**

You can create capacity and relationship reports that include the following information:

**Cluster Name**

The name that was assigned to the cluster when it was created.

**Cluster Type**

The type of the cluster. The cluster can be a Microsoft Cluster Server or a High-Availability Cluster Multiprocessing (HACMP) server.

**Cluster IP Address**

The IP address of the resource.

**Cluster Domain Name**

The domain name of the cluster.

**Data for servers in capacity and relationship reports**

You can include general information, capacity data, properties of components, and other information about servers in capacity and relationship reports.

**Information about servers**

You can create capacity and relationship reports that include the following information:

**Server Name**

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

**Server Domain Name**

The domain name of the server.

If a Tivoli Storage Productivity Center server uses one DNS server, and another server, which has a Storage Resource Agent installed, uses a different DNS server, the Tivoli Storage Productivity Center server might be displayed in some web-based GUI pages with different fully qualified host names. This is because the Tivoli Storage Productivity Center server is identified by two different DNS servers.

**Server Short Name**

The host name from the fully qualified domain name of the server. For example, if the fully qualified domain name of the server is `mycomputer.example.com`, the host name is `mycomputer`.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Server Swap Space (GiB)**

The amount of disk space on a resource that is available to store some of the contents of the RAM when the RAM is full.

**Server Total Disk Capacity (GiB)**

The amount of storage space that is on a server.

**Server Total Disk Available Space (GiB)**

The amount of unused storage space on all disks on a server.

**Server Total File System Available Space (GiB)**

The amount of unused storage space that is available on the file systems on a server.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Server Cluster Current Hosting Node**

The name of the server that is hosting the cluster resource group.

**Server Cluster Resource Type**

Shows whether the cluster is a node or a cluster resource group.

**Server IP Address**

The IP address of the resource.

**Server Is Virtual Machine**

Shows whether a resource is a virtual machine. If this value is Yes, the resource is a virtual machine.

**Server OS Type**

The operating system that is running on the server.

**Server OS Version**

The version number of the operating system that is running on the resource.

**Server Processor Architecture**

The architecture of the processor on the server or hypervisor. For example, the architecture of a processor might be Intel 64 bit (IA64) or Intel 32 bit (IA32).

**Server Processor Count**

The number of processors on the server or hypervisor.

**Server Processor Speed (MHz)**

The speed of the processor on the server or hypervisor.

**Server Processor Type**

Shows information about the processor, such as the family and model of the processor.

**Server RAM (GiB)**

The amount of RAM on the server or hypervisor.

**Server Time Zone**

The time zone in which a resource is located.

**Server Location**

The physical location of a server. The location is defined when a server is added to Tivoli Storage Productivity Center. You can add or edit the location of the server in the **Properties** pane of the server in the web-based GUI.

**Server Custom Tag 1, 2, and 3**

User-defined text that is associated with a server. You can add or edit the custom tags for a server in the **Properties** pane of the server in Tivoli Storage Productivity Center web-based GUI.

**Vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

**Server Model**

The model name or model number of the resource.

**Server Serial Number**

The serial number of the resource.

**Server Vendor**

The vendor who supplied the resource.

**Status information**

You can create capacity and relationship reports that include the following information:

**Server Status**

The status of a server, such as Normal, Warning, Error, Unreachable, Unknown, or Agentless. Use the status to determine the condition of the server, and if any actions must be taken. If this value is Agentless, the server is an unmanaged server.

**Server Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Server Last Data Collection**

The date and time when storage statistics were last collected from the resource.

**Server Last Start Time**

The last time that the resource was started.

**Data for file systems on servers in capacity and relationship reports**

You can include general information, capacity data, and status information about file systems on servers in capacity and relationship reports.

**Information about file systems on servers**

You can create capacity and relationship reports that include the following information:

**File System Mount Point**

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be /opt or /export/home.

**File System Export Name**

The name of the exported file system.

**File System Type**

The type of file system that the resource uses.

**Capacity and usage data****Tip:**

In the Tivoli Common Reporting interface, the capacity values of file systems on servers, on IBM SONAS systems, and on NAS systems are all displayed separately. The values are displayed in the following columns:

- Servers: File System Capacity (GiB)
- IBM SONAS systems and NAS systems: NAS File System Capacity (GiB)

However, in the web-based GUI, the capacity values of file systems are displayed in one column, Total File System Capacity (GiB). In the Tivoli Common Reporting interface, the File System Capacity (GiB) statistic for IBM SONAS systems and for NAS systems is blank.

You can create capacity and relationship reports that include the following information:

**File System Capacity (GiB)**

The amount of storage space on the file system of the resource.

**File System Directory Count**

The number of directories in the file systems that are on a server.

**File System File Count**

The number of files in the file system.

**File System Available Space (GiB)**

The amount of unused storage space in the file system of the resource.

**File System Available Inodes**

The number of unused inodes in file systems on the operating system.

**File System Used Inodes**

The number of used inodes in file systems on the operating system.

**Tip:** For Microsoft Windows systems, this property is blank.

**File System Maximum Files**

The maximum number of files that the file system on the resource can contain.

**File System Physical Size (GiB)**

The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

**File System Used Space (GiB)**

The amount of used storage space in the file system of the resource.

**Component properties**

You can create capacity and relationship reports that include the following information:

**File System Is Remote Mount**

Shows whether the file system is remotely mounted on another server. If this value is Yes, the file system is remotely mounted.

**Status information**

You can create capacity and relationship reports that include the following information:

**File System Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**File System Last Scan**

The date and time when file statistics were last collected from the resource.

**Data for logical volumes on servers in capacity and relationship reports**

You can include general information, capacity data, and status information about logical volumes on servers in capacity and relationship reports.

**Information about logical volumes on servers**

You can create capacity and relationship reports that include the following information:

**Logical Volume Path**

The path to a logical volume on a resource, for example /dev/hd1.

**Logical Volume Is Swap Space**

Shows whether a logical volume is used to store some of the contents of the RAM when the RAM is full. If this value is Yes, the logical volume is used to store some of the contents of the RAM.

**Logical Volume Mirror Count**

The number of mirrors that keep a synchronous copy of the resource.

**Logical Volume Type**

The type of server logical volume. For example, the logical volume type can be concatenation or stripe set.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Logical Volume Allocation on Server Disk (GiB)**

The amount of disk storage space that is allocated to the logical volume.

**Logical Volume Capacity (GiB)**

The amount of storage space that is on a logical volume.

**Logical Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Status information**

You can create capacity and relationship reports that include the following information:

**Logical Volume Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Data for volume groups on servers in capacity and relationship reports**

You can include general information and capacity data about volume groups on servers in capacity and relationship reports.

**Information about volume groups on servers**

You can create capacity and relationship reports that include the following information:

**Volume Group Name**

The name that was assigned to the volume group when it was added to the system.

**Volume Group Type**

The type of the volume manager that manages the volume group. The volume manager can be a logical volume manager (LVM) or a Veritas volume manager.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Volume Group Capacity (GiB)**

The amount of storage space that is on a volume group.

**Volume Group Available Space (GiB)**

The amount of unused storage space on a server volume group.

**Volume Group Number of Logical Volumes**

The number of logical volumes on a resource.



## Volume Group Number of Server Disks

The number of server disks in a volume group.

## Data for server disks in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about server disks in capacity and relationship reports.

### Information about server disks

You can create capacity and relationship reports that include the following information:

#### Server Disk Name

The path that the operating system uses for the server disk, for example `/dev/hdisk0`.

### Capacity and usage data

You can create capacity and relationship reports that include the following information:

#### Server Disk Capacity (GiB)

The amount of storage space that is on a server disk.

#### Server Disk Available Space (GiB)

The amount of unused storage space that is on a server disk.

### Component properties

You can create capacity and relationship reports that include the following information:

#### Server Disk Is Removable

Shows whether the disk can be removed from the server. If this value is Yes, the disk can be removed from the server.

#### Server Disk Multipathing Policy

The multipathing policy that is in effect for a disk. For example, the policy can be Round Robin, Load Balancing, Failover Only, or other policies.

### Status information

You can create capacity and relationship reports that include the following information:

#### Server Disk Status

The condition of the resource, for example normal, warning, or error.

#### Server Disk Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

#### Server Disk Firmware Version

The version number of the firmware that is running on the disk.

**Server Disk Model**

The model name or model number of the resource.

**Server Disk Serial Number**

The serial number of the resource.

**Server Disk Vendor**

The vendor who supplied the resource.

**Path information**

You can create capacity and relationship reports that include the following information:

**Server Disk Path Name**

The name of a path to a server disk. Typically the multipath driver defines this name.

**Server Disk Path Mode**

The online status of the path to the server disk.

**Server Disk Path State**

The state of the disk path shows whether the path is functioning correctly.

**Server Disk Path Status**

The aggregate status of the paths to the server disk.

**Server Disk Preferred Path**

The preferred path status of the path to the server disk, that is, whether the path is the preferred path for the disk. The multipath driver might use non-preferred paths if a failover occurs. The disk might have no preferred path, depending on the multipathing policy that is in effect for the disk.

**Server Disk LUN**

The logical unit within a small computer system interface target on a disk.

**Server Disk SCSI Target**

The bus address for the small computer system interface of the target on a disk.

**Server Disk Target WWPN**

The worldwide port name of the storage port that the volume is mapped to.

**Data for server groups in capacity and relationship reports**

You can include general information and properties of server groups in capacity and relationship reports.

**Information about server groups**

You can create capacity and relationship reports that include the following information:

**Server Group Name**

The name that was assigned to the server group when it was added to the system.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Server Group Creator**

The user name of the user who created the group.

**Server Group Last Modified**

The date and time that the resource was last modified.

**Server Group Last Modified by User**

The user name of the user who last modified the resource.

**Data for server controllers in capacity and relationship reports**

You can include general information about server controllers in capacity and relationship reports.

**Information about server controllers**

You can create capacity and relationship reports that include the following information:

**Server Controller Name**

The name of the disk controller on a server.

**Server Controller Status**

The status of a disk controller. Statuses include Normal, Warning, Error, and Unknown. Use the status to determine the condition of a controller, and if any actions must be taken. For example, if a controller has an Error status, take immediate action to correct the problem.

**Server Controller Description**

A short description of a disk controller that was provided by the manufacturer.

**Server Controller Type**

The type of disk controller, such as IDE, SCSI, Floppy, and RAID. Host Bus Adapters (HBAs) have a controller type of Fibre Channel Arbitrated Loop (FCAL).

**Server Controller Driver Version**

The version identifier of the device driver on the disk controller. This value is only available for HBAs.

**Server Controller Firmware**

The firmware version of the microcode on a disk controller. This value is only available for HBAs.

**Server Controller ROM Version**

The version of the read-only memory (ROM) on a controller. This value is only available for HBAs.

**Server Controller Hardware Version**

The hardware version identifier for a disk controller. This value is only available for HBAs.

**Server Controller WWN**

The worldwide name of a disk controller. This value is only available for HBAs.

**Server Controller Serial Number**

The serial number of the server controller.

**Server Controller Bus Address**

The SCSI bus address that is associated with a disk controller. This property is available only for SCSI disk controllers.

**Server Controller Bus Number**

The SCSI bus number that is associated with a disk controller. This property is available only for SCSI disk controllers.

**Data for multipath drivers on servers in capacity and relationship reports**

You can include name and version information about multipath drivers on servers in capacity and relationship reports.

**Information about multipath drivers on servers**

You can create capacity and relationship reports that include the following information:

**Server Multipath Driver Name**

The name of a multipath driver on a server.

**Server Multipath Driver Version**

The version number of a multipath driver on a server.

**Data for file system groups on servers in capacity and relationship reports**

You can include general information about file system groups on servers, and properties of file system groups in capacity and relationship reports.

**Information about file system groups on servers**

You can create capacity and relationship reports that include the following information:

**File System Group Name**

The name that was assigned to the file system group when it was added to the system.

**Component properties**

You can create capacity and relationship reports that include the following information:

**File System Group Creator**

The user name of the user who created the group.

**File System Group Last Modified**

The date and time that the resource was last modified.

**File System Group Last Modified by User**

The user name of the user who last modified the resource.

**Data for hypervisors in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about hypervisors in capacity and relationship reports.

## Information about hypervisors

You can create capacity and relationship reports that include the following information:

### Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

### Hypervisor Short Name

The host name from the fully qualified domain name of the hypervisor. For example, if the fully qualified domain name of the hypervisor is `mycomputer.example.com`, the host name is `mycomputer`.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### Hypervisor Kernel Memory (GiB)

The amount of memory on a hypervisor that is used for kernel memory. Kernel memory is used to perform internal operating system processes.

### Hypervisor Service Console Memory (GiB)

The amount of memory on a hypervisor that is used for the VMware ESX Service Console. Service Console is the operating system that is used to interact with ESX and the virtual machines that run on the hypervisor.

### Hypervisor Swap Space (GiB)

The amount of disk space on a resource that is available to store some of the contents of the RAM when the RAM is full.

### Hypervisor Total Disk Capacity (GiB)

The amount of storage space on disks on a hypervisor.

### Hypervisor Total Disk Available Space (GiB)

The amount of unused storage space on disks on the hypervisor.

### Hypervisor Total File System Available Space (GiB)

The amount of unused storage space that is available on the file systems on a hypervisor.

## Component properties

You can create capacity and relationship reports that include the following information:

### Hypervisor IP Address

The IP address of the resource.

### Hypervisor OS Type

The software that runs the hypervisor, that is, VMware ESX.

### Hypervisor OS Version

The version number of the operating system that is running on the resource.

### Hypervisor Processor Architecture

The architecture of the processor on the server or hypervisor. For example, the architecture of a processor might be Intel 64 bit (IA64) or Intel 32 bit (IA32).

**Hypervisor Processor Count**

The number of processors on the server or hypervisor.

**Hypervisor Processor Speed (MHz)**

The speed of the processor on the server or hypervisor.

**Hypervisor Processor Type**

Shows information about the processor, such as the family and model of the processor.

**Hypervisor RAM (GiB)**

The amount of RAM on the server or hypervisor.

**Hypervisor Software API Version**

The version number of the software API that runs on the hypervisor.

**Hypervisor Software Full Name**

The name, version number, and build number of the software that runs the hypervisor.

**Hypervisor Time Zone**

The time zone in which a resource is located.

**Hypervisor Location**

The physical location of a hypervisor. The location is defined when a hypervisor is added to Tivoli Storage Productivity Center. You can add or edit the location of the hypervisor in the **Properties** pane of the hypervisor in the web-based GUI.

**Hypervisor Custom Tag 1, 2, and 3**

User-defined text that is associated with a hypervisor. You can add or edit the custom tags for a hypervisor in the **Properties** pane of the hypervisor in the Tivoli Storage Productivity Center web-based GUI.

**Vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

**Hypervisor Model**

The model name or model number of the resource.

**Hypervisor Serial Number**

The serial number of the resource.

**Hypervisor Vendor**

The vendor who supplied the resource.

**Status information**

You can create capacity and relationship reports that include the following information:

**Hypervisor Status**

The condition of the resource, for example normal, warning, or error.

**Hypervisor Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Hypervisor Last Data Collection**

The date and time when storage statistics were last collected from the resource.

### **Hypervisor Last Start Time**

The last time that the resource was started.

## **Data for clusters on hypervisors in capacity and relationship reports**

You can include general information about clusters on hypervisors in capacity and relationship reports.

### **Information about clusters on hypervisors**

You can create capacity and relationship reports that include the following information:

#### **Hypervisor Cluster Name**

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

#### **Hypervisor Cluster Type**

The type of a cluster that is monitored in your storage environment. For example, the type might be VMware Cluster.

## **Data for file systems on hypervisors in capacity and relationship reports**

You can include general information, capacity data, and other information about file systems on hypervisors in capacity and relationship reports.

### **Information about file systems on hypervisors**

You can create capacity and relationship reports that include the following information:

#### **Hypervisor File System Mount Point**

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be /opt or /export/home.

#### **Hypervisor File System Export Name**

The name of the exported file system.

#### **Hypervisor File System Type**

The type of file system that the hypervisor uses.

### **Capacity and usage data**

You can create capacity and relationship reports that include the following information:

#### **Hypervisor File System Capacity (GiB)**

The amount of storage space on the file system of the resource.

#### **Hypervisor File System Directory Count**

The number of directories in the file systems that are on a server.

#### **Hypervisor File System File Count**

The number of files in the file system.

#### **Hypervisor File System Available Space (GiB)**

The amount of unused storage space in the file system of the resource.

**Hypervisor File System Available Inodes**

The number of unused inodes in file systems on the operating system.

**Hypervisor File System Used Inodes**

The number of used inodes in file systems on the operating system.

**Tip:** For Microsoft Windows systems, this property is blank.

**Hypervisor File System Maximum Files**

The maximum number of files that the file system on the resource can contain.

**Hypervisor File System Physical Size (GiB)**

The amount of physical storage space on the file systems on a resource.  
The physical size is the size of all the clusters that the file system uses.

**Hypervisor File System Used Space (GiB)**

The amount of used storage space in the file system of the resource.

**Status information**

You can create capacity and relationship reports that include the following information:

**Hypervisor File System Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Hypervisor File System Last Scan**

The date and time when file statistics were last collected from the resource.

**Data for data stores on hypervisors in capacity and relationship reports**

You can include general information, capacity data, and other information about data stores on hypervisors in capacity and relationship reports.

**Information about data stores on hypervisors**

You can create capacity and relationship reports that include the following information:

**Hypervisor Data Store**

The path to the VMware ESX data store on the hypervisor, for example `datastore_svc3c`.

**Hypervisor Data Store Cluster Name**

The name of the cluster, if the data store is a member of a data store cluster. If the data store is not a member of a data store cluster, no information is displayed.

**Hypervisor Data Store Type**

The type of file system that the hypervisor data store uses. For example, the type can be NFS or VMFS.

**Hypervisor NAS Name**

The name of the network-attached storage that the data store is obtained from.

**Hypervisor NAS Path Name**

The path name on the network-attached storage of the data store.



## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### **Hypervisor Data Store Capacity (GiB)**

The amount of storage space that is on the hypervisor data store.

### **Hypervisor Data Store Available Space (GiB)**

The amount of unused storage space on a hypervisor data store.

### **Hypervisor Data Store Allocation on Hypervisor Disk (GiB)**

The amount of storage space that is allocated to the hypervisor data store on hypervisor server disks.

### **Hypervisor Data Store VM File System Version**

The version number of the virtual machine file system on the hypervisor data store.

## Status information

You can create capacity and relationship reports that include the following information:

### **Hypervisor Data Store Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

## Data for server disks on hypervisors in capacity and relationship reports

You can include general information, capacity data, path information, and other information about server disks on hypervisors in capacity and relationship reports.

## Information about server disks

You can create capacity and relationship reports that include the following information:

### **Hypervisor Disk Name**

The path that the operating system uses for the server disk.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### **Hypervisor Disk Capacity (GiB)**

The amount of storage space that is on the hypervisor disk.

### **Hypervisor Disk Available Space (GiB)**

The amount of unused storage space on a hypervisor server disk.

## Component properties

You can create capacity and relationship reports that include the following information:

### **Hypervisor Disk Multipathing Policy**

The multipathing policy that is in effect for a disk. For example, the policy can be Round Robin, Load Balancing, Failover Only, or other policies.

## Status information

You can create capacity and relationship reports that include the following information:

### Hypervisor Disk Status

The condition of the resource, for example normal, warning, or error.

### Hypervisor Disk Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

## Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

### Hypervisor Disk Firmware Version

The version number of the firmware that is running on the disk.

### Hypervisor Disk Model

The model name or model number of the resource.

### Hypervisor Disk Serial Number

The serial number of the resource.

### Hypervisor Disk Vendor

The vendor who supplied the resource.

## Path information

You can create capacity and relationship reports that include the following information:

### Hypervisor Disk Path Name

The name of a path to a server disk. Typically the multipath driver defines this name.

### Hypervisor Disk Path State

The state of the disk path shows whether the path is functioning correctly.

### Hypervisor Disk Path Status

The aggregate status of the paths to the server disk.

### Hypervisor Disk Preferred Path

The preferred path status of the path to the server disk, that is, whether the path is the preferred path for the disk. The multipath driver might use non-preferred paths if a failover occurs. The disk might have no preferred path, depending on the multipathing policy that is in effect for the disk.

### Hypervisor Disk LUN

The logical unit within a small computer system interface target on a disk.

### Hypervisor Disk SCSI Target

The bus address for the small computer system interface of the target on a disk.

### Hypervisor Disk Target WWPN

The worldwide port name of the storage port that the volume is mapped to.

## Data for hypervisor controllers in capacity and relationship reports

You can include general information about hypervisor controllers in capacity and relationship reports.

### Information about hypervisor controllers

You can create capacity and relationship reports that include the following information:

#### Hypervisor Controller Name

The name of the disk controller on a hypervisor.

#### Hypervisor Controller Description

A short description of a disk controller that was provided by the manufacturer.

#### Hypervisor Controller Type

The type of disk controller, such as IDE, SCSI, Floppy, and RAID. Host Bus Adapters (HBAs) have a controller type of Fibre Channel Arbitrated Loop (FCAL).

#### Hypervisor Controller WWN

The worldwide name of a disk controller. This value is only available for HBAs.

#### Hypervisor Controller Bus Address

The SCSI bus address that is associated with a disk controller. This property is available only for SCSI disk controllers.

## Data for multipath drivers on hypervisors in capacity and relationship reports

You can include name and version information about multipath drivers on hypervisors in capacity and relationship reports.

### Information about multipath drivers on hypervisors

You can create capacity and relationship reports that include the following information:

#### Hypervisor Multipath Driver Name

The name of a multipath driver on a hypervisor.

#### Hypervisor Multipath Driver Version

The version number of a multipath driver on a hypervisor.

## Data for virtual machines managed by hypervisors in capacity and relationship reports

You can include general information, and file and disk information for virtual machines in capacity and relationship reports.

### Information about virtual machines

You can create capacity and relationship reports that include the following information:

#### Virtual Machine Name

The name that was assigned to the virtual machine when it was added to the system.

**Virtual Machine Configuration File**

The path to the virtual machine configuration file for a hypervisor virtual machine, for example [datastore1] example/example.vmx.

**Virtual Machine Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Virtual Machine Number of CPU**

The number of processors on the virtual machine.

**Virtual Machine Status**

The condition of the resource, for example normal, warning, or error.

**Virtual Machine Total Memory (GiB)**

The amount of memory on a hypervisor virtual machine.

**Virtual machine file information**

You can create capacity and relationship reports that include the following information:

**Virtual Machine File Path**

The path to the virtual machine configuration file, or to the virtual machine disk file for the virtual machine. The virtual machine configuration file is a .vmx file, and the virtual machine disk file is a .vmdk file.

**Virtual Machine File Logical Size (GiB)**

The logical size of a virtual machine disk file for a hypervisor virtual machine.

**Virtual Machine File Type**

The type of the file for the virtual machine. For example, the file can be a virtual machine configuration file, a virtual machine disk file, an ISO image file, or other types of virtual machine file.

**Virtual machine disk information**

You can create capacity and relationship reports that include the following information:

**Virtual Machine Disk Name**

The name that was assigned to the virtual machine disk when it was added to the system.

**Virtual Machine Disk SCSI Target**

The bus address for the small computer system interface of the target on a disk.

**Virtual Machine Disk Type**

The type of virtual machine disk, for example, collection mirroring or logical disk mirroring.

**Data for network-attached storage systems in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about network-attached storage (NAS) systems in capacity and relationship reports.

## Information about NAS systems

You can create capacity and relationship reports that include the following information:

### NAS Name

The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

### NAS Domain Name

The domain name of the cluster, server, or NAS server.

### NAS Short Name

The host name from the fully qualified domain name of the NAS device. For example, if the fully qualified domain name of the NAS device is `mycomputer.example.com`, the host name is `mycomputer`.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### NAS Total Disk Capacity (GiB)

The amount of storage space on all disks on a NAS resource.

### NAS Total Disk Available Space (GiB)

The amount of unused storage space on all disks on a NAS resource.

### NAS Total File System Available Space (GiB)

The amount of unused storage space on the file systems on a NAS resource.

## Component properties

You can create capacity and relationship reports that include the following information:

### NAS OS Type

The operating system that is running on the server.

### NAS OS Version

The version number of the operating system that is running on the resource.

### NAS Location

The physical location of a NAS system that is probed as a file server. You can add or edit the location of the NAS system on the **File Storage** tab in the **Storage Systems** pane in the web-based GUI.

### NAS Custom Tag 1, 2, and 3

User-defined text that is associated with a NAS system that is probed as a file server. You can add or edit the custom tags for a NAS system on the **File Storage** tab in the **Storage Systems** pane in the Tivoli Storage Productivity Center web-based GUI.

## Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

**NAS Model**

The model name or model number of the resource.

**NAS Vendor**

The vendor who supplied the resource.

**NAS Serial Number**

The serial number of the resource.

**Status information**

You can create capacity and relationship reports that include the following information:

**NAS Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**NAS Last Data Collection**

The date and time when storage statistics were last collected from the resource.

**NAS Last Start Time**

The last time that the resource was started.

**NAS system components information**

You can create capacity and relationship reports that include the following information:

**NAS System Name**

The name that was assigned to the NAS system when it was added to the system.

**NAS System IP Address**

The IP address of the resource.

**NAS System Location**

The physical location of a NAS system that is probed as a storage system. The location is defined when a NAS system is added to Tivoli Storage Productivity Center. You can add or edit the location of the NAS system on the **File Storage** tab in the **Storage Systems** pane in the web-based GUI.

**NAS System Custom Tag 1, 2, and 3**

User-defined text that is associated with a NAS system that is probed as a storage system. You can add or edit the custom tags for a NAS system on the **File Storage** tab in the **Storage Systems** pane in the Tivoli Storage Productivity Center web-based GUI.

**Data for file systems on network-attached storage systems in capacity and relationship reports**

You can include general information, capacity data, and other information about file systems on network-attached storage (NAS) systems in capacity and relationship reports.

**Information about file systems on NAS systems**

You can create capacity and relationship reports that include the following information:

**NAS File System Mount Point**

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be /opt or /export/home.

**NAS File System Export Name**

The name of the exported file system.

**NAS File System Type**

The type of file system that the resource uses.

**Capacity and usage data****Tip:**

In the Tivoli Common Reporting interface, the capacity values of file systems on servers, on IBM SONAS systems, and on NAS systems are all displayed separately. The values are displayed in the following columns:

- Servers: File System Capacity (GiB)
- IBM SONAS systems and NAS systems: NAS File System Capacity (GiB)

However, in the web-based GUI, the capacity values of file systems are displayed in one column, Total File System Capacity (GiB). In the Tivoli Common Reporting interface, the File System Capacity (GiB) statistic for IBM SONAS systems and for NAS systems is blank.

You can create capacity and relationship reports that include the following information:

**NAS File System Capacity (GiB)**

The amount of storage space on the file system of the resource.

**NAS File System Directory Count**

The number of directories in the file systems that are on a server.

**NAS File System File Count**

The number of files in the file system.

**NAS File System Available Space (GiB)**

The amount of unused storage space in the file system of the resource.

**NAS File System Available Inodes**

The number of unused inodes on the file system on a NAS server.

**NAS File System Used Inodes**

The number of used inodes on the file system on a NAS server.

**NAS File System Maximum Files**

The maximum number of files that the file system on the resource can contain.

**NAS File System Physical Size (GiB)**

The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

**NAS File System Used Space (GiB)**

The amount of used storage space in the file system of the resource.

## Component properties

You can create capacity and relationship reports that include the following information:

### NAS File System Custom Tag 1, 2, and 3

User-defined text that is associated with a NAS file system. You can add or edit the custom tags for a NAS file system in the **Properties** pane of the resource in the Tivoli Storage Productivity Center web-based GUI.

## Status information

You can create capacity and relationship reports that include the following information:

### NAS File System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### NAS File System Last Scan

The date and time when file statistics were last collected from the resource.

## Data for network-attached storage exports in capacity and relationship reports

You can include general information and status information about network-attached storage (NAS) exports in capacity and relationship reports.

## Information about network-attached storage exports

You can create capacity and relationship reports that include the following information:

### NAS Export Name

The name of the network-attached storage export. Use this name to mount the network-attached storage export on a server.

### NAS Export Path

The path to a logical volume on a resource, for example `/dev/hd1`.

### NAS Export State

Shows whether the network-attached storage export is active.

## Status information

You can create capacity and relationship reports that include the following information:

### NAS Export Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

## Data for logical volumes on network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, and other information about logical volumes on network-attached storage (NAS) systems in capacity and relationship reports.



## Information about logical volumes on NAS systems

You can create capacity and relationship reports that include the following information:

### NAS Logical Volume Name

The name that was assigned to the NAS logical volume when it was added to the system.

### NAS Logical Volume Type

The type of NAS logical volume. The logical volume type can be a NetApp traditional volume or a NetApp flexible volume.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### NAS Logical Volume Capacity (GiB)

The amount of storage space on the NAS logical volume.

### NAS Logical Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

## Status information

You can create capacity and relationship reports that include the following information:

### NAS Logical Volume Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### NAS Logical Volume Last Scan

The date and time when file statistics were last collected from the resource.

## Data for network-attached storage disks in capacity and relationship reports

You can include general information, capacity data, properties, and other information about network-attached storage (NAS) disks in capacity and relationship reports.

## Information about NAS disks

You can create capacity and relationship reports that include the following information:

### NAS Disk Name

The name that was assigned to the NAS disk when it was added to the system.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### NAS Disk Total Capacity (GiB)

The amount of storage space that is on the NAS disk.

### **NAS Disk Available Space (GiB)**

The amount of unused storage space on a NAS disk.

### **Status information**

You can create capacity and relationship reports that include the following information:

#### **NAS Disk Status**

The condition of the resource, for example normal, warning, or error.

#### **NAS Disk Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### **IBM Scale Out Network Attached Storage and Storwize V7000 Unified information**

You can create capacity and relationship reports that include the following information:

#### **NSD Available Fragments**

The number of fragments of disk space that are unused on the Network Shared Disk (NSD).

#### **NSD Failover Group**

The failover group that the NSD belongs to. A failover group is a group of resources that are grouped together to ensure that some of those resources are available if one or more of the resources fail.

#### **NSD Type**

The type of the NSD. For example, the type can be:

- Data and metadata
- Metadata only
- Data only

#### **NSD Last Data Collection**

The date and time when storage statistics were last collected from the resource.

#### **NSD Custom Tag 1, 2, and 3**

User-defined text that is associated with a Network Shared Disk (NSD).

You can add or edit the custom tags for an NSD in the **Properties** pane of the resource in the Tivoli Storage Productivity Center web-based GUI.

### **NetApp information**

You can create capacity and relationship reports that include the following information:

#### **NAS Disk Flag**

Shows whether the NAS disk is a spare device. If this value is Spare device, the NAS disk is a spare device.

#### **NAS Disk Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

**NAS Disk Is Removable**

Shows whether the disk can be removed from the server. If this value is Yes, the disk can be removed from the server.

**NAS Disk Is Solid State**

Shows whether the disk is a solid-state drive.

**NetApp paths information**

You can create capacity and relationship reports that include the following information:

**NAS Disk LUN**

The logical unit within a small computer system interface target on a disk.

**NAS Disk SCSI Target**

The bus address for the small computer system interface of the target on a disk.

**NAS Disk Target WWPN**

The worldwide port name of the storage port that the volume is mapped to.

**NetApp vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

**NAS Disk Firmware Version**

The version number of the firmware that is running on the disk.

**NAS Disk Model**

The model name or model number of the resource.

**NAS Disk Serial Number**

The serial number of the resource.

**NAS Disk Vendor**

The vendor who supplied the resource.

**Data for pools on IBM SONAS and Storwize V7000 Unified systems in capacity and relationship reports**

You can include general information and capacity data about pools on IBM Scale Out Network Attached Storage and Storwize V7000 Unified systems in capacity and relationship reports.

**Information about pools on network-attached storage systems**

You can create capacity and relationship reports that include the following information:

**NAS Pool Name**

The name that was assigned to the NAS pool when it was added to the system.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**NAS Pool Capacity (GiB)**

The amount of storage space that is in a NAS pool.

**NAS Pool Available Space (GiB)**

The amount of unused storage space in a NAS pool.

**NAS Pool Number of Logical Volumes**

The number of logical volumes on a resource.

**Data for IBM SONAS and Storwize V7000 Unified filesets in capacity and relationship reports**

You can include general information, capacity data, and other information about IBM Scale Out Network Attached Storage and Storwize V7000 Unified filesets in capacity and relationship reports.

**Information about network-attached storage filesets**

You can create capacity and relationship reports that include the following information:

**NAS Fileset Name**

The name that was assigned to the network-attached storage (NAS) fileset when it was added to the system.

**NAS Fileset Comment**

A comment that is associated with the NAS fileset.

**NAS Fileset Path**

The path to a logical volume on a resource, for example `/dev/hd1`.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**NAS Fileset Used Space (GiB)**

The amount of used storage space in the fileset of the NAS resource.

**NAS Fileset Used Inodes**

The number of used inodes in file systems on operating systems such as AIX, Linux, Solaris, or HP-UX.

**Status information**

You can create capacity and relationship reports that include the following information:

**NAS Fileset Status**

Shows whether the network-attached storage fileset is linked or unlinked.

**NAS Fileset Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Data for storage virtualizers in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about storage virtualizers in capacity and relationship reports.

## Information about storage virtualizers

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Total Pool Capacity (GiB)

The total amount of storage space in pools on the storage virtualizer.

### Storage Virtualizer Pool Used Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

### Storage Virtualizer Pool Available Space (GiB)

The amount of space in a pool that is not reserved for volumes.

### Storage Virtualizer Total Volume Capacity (GiB)

The amount of virtual volume space on all volumes on the storage virtualizer.

### Storage Virtualizer Volume Allocated Space (GiB)

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

**Storage Virtualizer Volume Assigned Space (GiB)**

The amount of virtual volume space that is assigned to the server or storage virtualizer.

**Storage Virtualizer MDisk Available Space (GiB)**

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

**Storage Virtualizer MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage Virtualizer Disk Capacity (GiB)**

The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

**Storage Virtualizer Number of MDisks**

The number of managed disks on a resource.

**Storage Virtualizer Number of Volumes**

The number of volumes on a resource.

**Storage Virtualizer Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Virtualizer Tier Available Space (GiB)**

The amount of unused tier storage of a particular type. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Tier Capacity (GiB)**

The amount of tier storage of a particular type. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Physical Allocation Percentage**

The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Virtualizer Virtual Allocation Percentage**

The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Shortfall Percentage**

The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

### **Component properties**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Access Information**

The URL from which you can access the storage system or storage virtualizer.

**Storage Virtualizer Alias**

Shows an alternative name for a storage virtualizer.

**Storage Virtualizer Code Level**

The microcode or firmware level of the storage virtualizer.

**Storage Virtualizer Time Zone**

The time zone in which a resource is located.

**Storage Virtualizer Location**

The physical location of a storage virtualizer. The location is defined when a storage virtualizer is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage virtualizer in the **Properties** pane for the storage system in the web-based GUI.

**Storage Virtualizer Custom Tag 1, 2, and 3**

User-defined text that is associated with a storage virtualizer. You can add or edit the custom tags for a storage virtualizer on the **Properties** page of the storage system in the Tivoli Storage Productivity Center web-based GUI.

**Storage Virtualizer User Provided Name**

The name that was specified for the storage virtualizer in the storage environment.

**Storage Virtualizer Tier Type**

The type of tier that is used on the storage virtualizer. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD). The tier is configured to be controlled by the Easy Tier function.

This value is available only for SAN Volume Controller and Storwize V7000 storage systems.

**Storage Virtualizer Is Compression Active**

Shows whether the compression feature is enabled on the storage virtualizer. If this value is Yes, the compression feature is enabled.

**Storage Virtualizer Configuration**

Shows that a storage virtualizer is configured to virtualize back-end storage. This value is always **Storage Virtualizer**.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Status**

The condition of the resource, for example normal, warning, or error.

**Storage Virtualizer Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Virtualizer Last Data Collection**

The date and time when storage statistics were last collected from the resource.



### **Storage Virtualizer Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

### **Vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer Type**

The type of storage virtualizer. For example, the storage virtualizer can be an **IBM Storwize V7000** system, an **IBM SAN Volume Controller** system, or another type of storage virtualizer.

#### **Storage Virtualizer Vendor**

The vendor who supplied the resource.

#### **Storage Virtualizer Model**

The model name or model number of the storage virtualizer.

#### **Storage Virtualizer Machine Type**

The machine type of a storage virtualizer. For example, the storage virtualizer can be **Storwize V7000 - 2076**, **SAN Volume Controller - 2145**, or another type of storage virtualizer.

#### **Storage Virtualizer IP Address**

The IP address of the resource.

#### **Storage Virtualizer Serial Number**

The serial number of the resource.

### **Data for storage virtualizer pools in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about storage virtualizer pools in capacity and relationship reports.

### **Information about storage virtualizer pools**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer Pool Name**

The name that was assigned to the pool when it was added to the system.

### **Capacity and usage data**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer Pool Capacity (GiB)**

The total amount of storage space in a pool.

#### **Storage Virtualizer Pool Available Space (GiB)**

The amount of unused storage space in a storage virtualizer pool.

#### **Storage Virtualizer Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Virtualizer Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

**Storage Virtualizer Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Virtualizer Pool Extent Size (MiB)**

The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Virtualizer Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Virtualizer Pool Tier Capacity (GiB)**

The amount of pool tier storage of a particular type. The type can be hard disk drive (HDD) or solid-state drive (SSD) and is identified in the Storage Virtualizer Pool Tier Type property. The tier is configured to be controlled by the Easy Tier function.

This value is available only for SAN Volume Controller and Storwize V7000 storage systems.

**Storage Virtualizer Pool Tier Available Space (GiB)**

The amount of pool tier storage of a particular type that is unused. The type can be hard disk drive (HDD) or solid-state drive (SSD) and is identified in the Storage Virtualizer Pool Tier Type property. The tier is configured to be controlled by the Easy Tier function.

This value is available only for SAN Volume Controller and Storwize V7000 storage systems.

**Storage Virtualizer Pool Tier Number of MDisks**

The number of managed disks that is on a storage virtualizer pool. The tier is configured to be controlled by the Easy Tier function.

This value is available only for SAN Volume Controller and Storwize V7000 storage systems.

**Storage Virtualizer Pool Number of Storage Volumes**

The number of volumes in the storage virtualizer pool.

**Storage Virtualizer Pool Number of MDisks**

The number of managed disks on a resource.

**Storage Virtualizer Pool Maximum I/O Capability**

The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Virtualizer Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Virtualizer Pool Back-End Storage System Type
- Storage Virtualizer Pool Back-End Storage RAID Level
- Storage Virtualizer Pool Back-End Storage Disk Type

This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Virtualizer Pool Back-End Storage Disks**

The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Virtualizer Pool Physical Allocation Percentage**

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Virtualizer Pool Virtual Allocation Percentage**

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

#### **Storage Virtualizer Pool Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{total volume capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

#### **Storage Virtualizer Pool Unallocated Volume Space (GiB)**

The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$\text{total volume capacity} - \text{allocated space}$$

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

#### **Storage Virtualizer Pool Unused Volume Space (GiB)**

The amount of space that is allocated to the volumes in a pool and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{allocated space} - \text{used pool space}$$

**Storage Virtualizer Pool Unallocatable Volume Space (GiB)**

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

**Storage Virtualizer Pool Used Space Percentage**

The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

**Storage Virtualizer Pool Compressed Virtual Capacity (GiB)**

The total virtual capacity of all the volumes that are compressed in a storage virtualizer pool.

**Storage Virtualizer Pool Compressed Capacity (GiB)**

The amount of storage space that is used by compressed volumes in a storage virtualizer pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

**Storage Virtualizer Pool Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Pool Is Encryptable**

Shows whether the storage virtualizer pool can be encrypted. If this value is Yes, the storage virtualizer pool can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool Is Encrypted**

Shows whether the storage virtualizer pool is encrypted. If this value is Yes, the storage virtualizer pool is encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool Is Solid State**

Shows whether there are solid-state drives in the pool. This property can contain the following values:

**Mixed** The pool contains both hard disk drives and solid-state drives.

**Non solid state**

The pool contains no solid-state drives.

**Solid state**

The pool contains a least one solid-state drive.

**Storage Virtualizer Pool Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Virtualizer Pool RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Virtualizer Pool Warning Level**

The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Pool Tier**

The tier level of pools on storage virtualizers. For example, the tier can be 1, 2, 3, or another tier number.

**Storage Virtualizer Pool Tier Type**

The type of tier that is used on the storage virtualizer pool. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD). The tier is configured to be controlled by the Easy Tier function.

This value is available only for SAN Volume Controller and Storwize V7000 storage systems.

**Storage Virtualizer Pool Is Easy Tier**

Shows whether the IBM System Storage Easy Tier feature is enabled. The IBM System Storage Easy Tier feature can be set to auto, on, or off.

The IBM System Storage Easy Tier feature automatically helps optimize solid-state storage deployments in multitier systems.

**Storage Virtualizer Pool Back-End Storage System Type**

The type of storage system that is providing storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Virtualizer Pool Back-End Storage RAID Level**

The RAID level of the volumes on the back-end storage system that are providing storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Virtualizer Pool Back-End Storage Disk Type**

The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Virtualizer Pool Custom Tag 1, 2, and 3**

User-defined text that is associated with a storage virtualizer pool. You can add or edit the custom tags for a storage virtualizer pool on the **Properties** notebook of the pool in the Tivoli Storage Productivity Center web-based GUI.

**Storage Virtualizer Pool Is Compression Active**

Shows whether the compression feature is enabled on the storage virtualizer pool. If this value is Yes, the compression feature is enabled.

## Status information

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Pool Status

The condition of the resource, for example normal, warning, or error.

### Storage Virtualizer Pool Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

### Storage Virtualizer Pool Easy Tier Status

Shows whether the Easy Tier function is active or inactive. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

## Data for storage virtualizer volumes in capacity and relationship reports

You can include general information, capacity data, properties, information about copies of volumes, and other information about storage virtualizer volumes in capacity and relationship reports.

## Information about storage virtualizer volumes

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

### Storage Virtualizer Volume Assigned Host Connection

The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

### Storage Virtualizer Volume Host Connection OS Type

The operating system type of the server or hypervisor that the virtualizer volume is assigned to.

### Storage Virtualizer Volume Preferred Node Name

The number of the storage virtualizer node that is used for I/O operations for the storage virtualizer volume.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

### Storage Virtualizer Volume Allocated Space (GiB)

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

#### **Storage Virtualizer Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

#### **Storage Virtualizer Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

#### **Storage Virtualizer Volume Used Allocated Space Percentage**

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 volumes.

#### **Storage Virtualizer Volume Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine the value:

$$\text{capacity} - \text{allocated space}$$

#### **Storage Virtualizer Volume Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

#### **Storage Virtualizer Volume Unused Space (GiB)**

The amount of space that is allocated to a volume and is not yet used.



Tivoli Storage Productivity Center uses the following formula to determine this value:

allocated space - used space

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 pools.

#### **Storage Virtualizer Volume Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$(\text{unallocatable volume space} \div \text{unallocated volume space}) \times 100$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$(100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space}) \div 100 \text{ GiB unallocated volume space} \times 100$

This value is only available for pools with thin-provisioned volumes.

#### **Storage Virtualizer Volume Unallocatable Space (GiB)**

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

#### **Storage Virtualizer Volume Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

#### **Storage Virtualizer Volume Tier Capacity (GiB)**

The amount of tier storage of a particular type on a virtualizer volume. The type is identified in the Storage Virtualizer Volume Tier Type property. The tier is configured to be controlled by the Easy Tier function.

#### **Storage Virtualizer Volume Tier Available Space (GiB)**

The amount of tier storage of a particular type that is unused on a virtualizer volume. The type is identified in the Storage Virtualizer Volume Tier Type property. The tier is configured to be controlled by the Easy Tier function.

## Component properties

You can create capacity and relationship reports that include the following information:

### **Storage Virtualizer Volume Type**

The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Volume WWN**

The worldwide name of the volume.

### **Storage Virtualizer Volume Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

### **Storage Virtualizer Volume Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Storage Virtualizer Volume Is Assigned**

Shows whether the volume is assigned to a server. If this value is Yes, the volume is assigned to a server.

### **Storage Virtualizer Volume Is Auto Expand Enabled**

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is Yes, the volume or volume copy can allocate new extents from a pool automatically.

### **Storage Virtualizer Volume Is Encryptable**

Shows whether the resource can be encrypted. If this value is Yes, the resource can be encrypted.

### **Storage Virtualizer Volume Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

### **Storage Virtualizer Volume Is Formatted**

Shows whether a storage virtualizer volume is formatted. If this value is Yes, the storage virtualizer volume is formatted.

### **Storage Virtualizer Volume Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

### **Storage Virtualizer Volume Metro Mirror Name**

The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

### **Storage Virtualizer Volume Mirror Count**

The number of mirrors that keep a synchronous copy of the resource.

**Storage Virtualizer Volume Throttle**

The maximum number of commands that the storage virtualizer volume can queue.

**Storage Virtualizer Volume Warning Level**

The percentage of used capacity of the resource at which a warning is generated.

**Storage Virtualizer Volume FlashCopy Relationship**

Shows whether a volume on a storage virtualizer is in a FlashCopy relationship. This property can contain the following values:

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Not in a FlashCopy Relationship**

The volume is not in a FlashCopy relationship.

**Storage Virtualizer Volume Remote Copy Relationship**

Shows whether a volume on a storage virtualizer is in a remote copy relationship. This property can contain the following values:

**Not in a Copy Relationship**

The volume is not in a remote copy relationship.

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Storage Virtualizer Volume RAID Level**

The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Virtualizer Volume Is Compressed**

Shows whether the storage virtualizer volume is compressed. If this value is Yes, the volume is compressed.

**Storage Virtualizer Volume Is Easy Tier**

Shows how the Easy Tier function is enabled on a virtualizer volume. For example, if the value of this column is auto, the Easy Tier function is enabled in automatic mode.

**Storage Virtualizer Volume Tier Type**

The type of tier that is used on the virtualizer volume. The tier is configured to be controlled by the Easy Tier function.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Volume Status**

The condition of the resource, for example normal, warning, or error.

**Storage Virtualizer Volume Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Virtualizer Volume Native Status**

Shows the level of access of nodes in the system to the storage virtualizer managed disk or volume. The level of access can be online, offline, degraded, or excluded.

**Storage Virtualizer Volume Easy Tier Status**

Shows whether the Easy Tier function is active or inactive on the virtualizer volume.

**Volume copy information**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Volume Copy ID**

The identifier of a volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Volume Copy Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

**Storage Virtualizer Volume Copy Is Auto Expand Enabled**

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is Yes, the volume or volume copy can allocate new extents from a pool automatically.

**Storage Virtualizer Volume Copy Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Virtualizer Volume Copy Is Primary**

Shows whether the volume copy is the primary volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Volume Copy Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Virtualizer Volume Copy Total Number of Extents**

The number of extents that are used by the storage virtualizer volume copy.

**Storage Virtualizer Volume Copy Type**

The type of copy of the storage virtualizer volume. For example, the copy of the storage virtualizer volume can be sequential, a mirror copy, or another type of copy.

**Storage Virtualizer Volume Copy Warning Level**

The percentage of used capacity of the resource at which a warning is generated.

**Storage Virtualizer Volume Copy Used Allocated Space Percentage**

The percentage of reserved space for a volume copy that is being used. This value is always less than or equal to 100% because a volume copy cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{used space} \div \text{allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 pools.

**Storage Virtualizer Volume Copy Used Space (GiB)**

The amount of allocated space that is used by a volume copy.

**Storage Virtualizer Volume Copy Allocated Space (GiB)**

The amount of space that is reserved for a volume copy.

The space that is allocated for a thin-provisioned volume copy is less than its virtual capacity, which is shown in Storage Virtualizer Capacity (GiB).

This value is equal to Storage Virtualizer Volume Copy Used Space (GiB) for volume copies that are not thin provisioned.

**Storage Virtualizer Volume Copy Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume copy.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{capacity} - \text{allocated space}$$

**Storage Virtualizer Volume Copy Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume copy. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{allocated Space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volume copies is 50 GiB.

**Storage Virtualizer Volume Copy Unused Space (GiB)**

The amount of space that is allocated to a volume copy and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 pools.

**Storage Virtualizer Volume Copy Shortfall Percentage**

The percentage of the remaining unallocated volume copy space in a pool that is not available to be allocated to a volume copy.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{unallocatable space} \div (\text{capacity} - \text{real capacity})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$$((100 \text{ GiB unallocatable volume copy space} - 70 \text{ GiB remaining pool space}) \div 100 \text{ GiB unallocated volume space}) \times 100$$

Available only for thin-provisioned volume copies.

#### **Storage Virtualizer Volume Copy Unallocatable Space (GiB)**

The unallocatable storage space in GiB of a thin-provisioned volume copy. Unallocatable space cannot be supplied by the pool to which the volume copy belongs.

#### **Storage Virtualizer Volume Copy Is Compressed**

Shows whether the volume copy is compressed. If this value is Yes, the volume copy is compressed.

#### **Storage Virtualizer Volume Copy Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed space on the volume copy is uncompressed. For example, if 100 GiB of data is compressed, the value is 100.

### **Data for managed disks on storage virtualizers in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about managed disks on storage virtualizers in capacity and relationship reports.

#### **Information about managed disks on storage virtualizers**

You can create capacity and relationship reports that include the following information:

##### **Storage Virtualizer MDisk Name**

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

##### **Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage Virtualizer MDisk Available Space (GiB)**

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

**Storage Virtualizer MDisk Strip Size (KB)**

The RAID strip size on a managed disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer MDisk RAID Level**

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer MDisk Spare Goal**

The number of spare drives that are required to maintain redundancy. Use spare drives to protect against drive failures in the array on a managed disk on a storage virtualizer.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer MDisk Type**

The type of managed disk on a storage virtualizer. For example, the storage virtualizer can be a local or virtual managed disk.

**Storage Virtualizer MDisk Mode**

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

**Array** Extents are provided from local disks.

**Managed**

Extents are provided from a back-end storage volume.

**Unmanaged**

The managed disk is not used in the system.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer MDisk Is Balanced**

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is Yes, the LUNs are balanced.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer MDisk Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted. This value is

not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

#### **Storage Virtualizer MDisk Write Verify**

Shows whether all write operations on a managed disk on a storage virtualizer are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is Yes, all write operations are verified by a follow-up read operation.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Status information**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer MDisk Status**

The condition of the resource, for example normal, warning, or error.

#### **Storage Virtualizer MDisk Native Status**

Shows the level of access of nodes in the system to a managed disk or volume on a storage virtualizer. The level of access can be online, offline, degraded, or excluded.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

### **Data for local disks on storage virtualizers in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about local disks on storage virtualizers in capacity and relationship reports.

### **Information about local disks on storage virtualizers**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer Local Disk Name**

The name that was assigned to the local disk on the storage virtualizer when it was added to the system.

### **Capacity and usage data**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer Local Disk Capacity (GiB)**

The amount of storage space on the local disk.

### **Component properties**

You can create capacity and relationship reports that include the following information:

#### **Storage Virtualizer Local Disk Is Encryptable**

Shows whether the local disk can be encrypted. If this value is Yes, the local disk can be encrypted.



This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Local Disk Is Encrypted**

Shows whether the local disk is encrypted. If this value is Yes, the local disk is encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Local Disk Storage Class**

The storage technology of the local disk on a storage virtualizer. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD). This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Local Disk Is Solid State**

Shows whether the disk is a solid-state drive. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Local Disk Speed (RPM)**

The speed of the local disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Local Disk Tag**

A number that identifies the array to which a local disk belongs.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Local Disk Status**

The condition of the resource, for example normal, warning, or error.

**Vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

**Storage Virtualizer Local Disk Firmware Version**

The version number of the firmware that is running on the local disk. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

**Storage Virtualizer Local Disk Model**

The model name or model number of the resource.

**Storage Virtualizer Local Disk Serial Number**

The serial number of the resource.

**Storage Virtualizer Local Disk Vendor**

The vendor who supplied the resource.

**Data for storage virtualizer groups in capacity and relationship reports**

You can include the name of storage virtualizer groups and properties for storage virtualizer groups in capacity and relationship reports.

## Information about storage virtualizer groups

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Group Name

A user-defined name of the storage virtualizer group.

## Component properties

You can create capacity and relationship reports that include the following information:

### Storage Virtualizer Group Creator

The user name of the user who created the storage virtualizer group.

### Storage Virtualizer Group Last Modified

The date and time that the storage virtualizer group was last modified.

### Storage Virtualizer Group Last Modified by User

The user name of the user who last modified the storage virtualizer group.

## Data for storage systems in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about storage systems in capacity and relationship reports.

## Information about storage systems

You can create capacity and relationship reports that include the following information:

### Storage System Name

A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

## Capacity and usage data

You can create capacity and relationship reports that include the following information:

### Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

### Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

### Storage System Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

### Storage System Volume Capacity (GiB)

The amount of space on all volumes on the storage system.

### Storage System Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage System Assigned Volume Space (GiB)**

The amount of space on the storage system that is on volumes that are assigned to a server or storage virtualizer.

**Storage System Unallocated Disk Space (GiB)**

The amount of disk space that can be added to a pool.

**Storage System Number of Disks**

The number of disks on the storage system.

**Storage System Number of Volumes**

The number of volumes on a resource.

**Storage System Cache (GiB)**

The size of the cache on the storage system. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Real Available Pool Space (GiB)**

The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Real Configured Pool Space (GiB)**

The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage System Volume Capacity for z/OS (GiB)**

The amount of space on all volumes on the storage system that the z/OS operating system can use.

**Storage System Volume Capacity Assigned to MDisks (GiB)**

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

**Storage System Disk Capacity (GiB)**

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System Total Disk Capacity (GiB)**

The amount of space on physical disks on a storage system, including spare disks.

**Storage System Physical Allocation Percentage**

The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{allocated pool space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage System Virtual Allocation Percentage**

The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

### **Storage System Shortfall Percentage**

The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB  $\div$  100 GiB  $\times$  100).

This value is only available for pools with thin-provisioned volumes.

### **Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### **Storage System Tier Available Space (GiB)**

The amount of unused space in a tier. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### **Storage System Tier Capacity (GiB)**

The total amount of storage space in a tier. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

## Component properties

You can create capacity and relationship reports that include the following information:

### Storage System Access Information

The URL from which you can access the storage system or storage virtualizer.

### Storage System Code Level

The Shared Ethernet Adapter level of a DS8000 storage system. For other storage systems, this value is the firmware version.

### Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

### Storage System Custom Tag 1, 2, and 3

User-defined text that is associated with a storage system. You can add or edit the custom tags for a storage system in the **Properties** pane of the storage system in the Tivoli Storage Productivity Center web-based GUI.

### Storage System User Provided Name

The name that was specified for the storage system in the storage environment.

### Storage System Tier Type

The type of tier that is used on a storage system. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD).

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

### Storage System Is Compression Active

Shows whether the compression feature is enabled on the storage system. If this value is **Yes**, the compression feature is enabled.

## Status information

You can create capacity and relationship reports that include the following information:

### Storage System Status

The condition of the resource, for example normal, warning, or error.

### Storage System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

### Storage System Last Data Collection

The date and time when storage statistics were last collected from the resource.

**Storage System Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

**Vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

**Storage System Type**

The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

**Storage System Vendor**

The vendor who supplied the resource.

**Storage System Model**

The model name or model number of the storage system.

**Storage System Machine Type**

The machine type of a storage system. For example, the storage system can be **DS8800**, **DS8700**, **XIV - 2812**, or another type of storage system.

**Storage System IP Address**

The IP address of the resource.

**Storage System Serial Number**

The serial number of the resource.

**Data for storage system pools in capacity and relationship reports**

You can include general information, capacity data, properties of components, and other information about storage system pools in capacity and relationship reports.

**Information about storage systems pools**

You can create capacity and relationship reports that include the following information:

**Storage Pool Name**

The name that was assigned to the pool when it was added to the system.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Storage Pool Capacity (GiB)**

The total amount of storage space in a pool.

**Storage Pool Available Space (GiB)**

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Tivoli Storage Productivity Center uses the following formula to determine this value:

pool capacity - used space

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

**Storage Pool Assigned Volume Space (GiB)**

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

**Storage Pool Unassigned Volume Space (GiB)**

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

**Storage Pool Total Volume Capacity (GiB)**

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

**Storage Pool Real Configured Space (GiB)**

The amount of storage space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage Pool Used Volume Space (GiB)**

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Pool Real Available Space (GiB)**

The amount of unused space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

**Storage Pool Repository Capacity (GiB)**

The amount of space on all extents in the repository of a pool. This space can be used to allocate track space-efficient volumes. This attribute applies only to DS8000 storage systems.

**Storage Pool Used Repository Space (GiB)**

The amount of space on all extents in the repository of a pool that are allocated. This space can be used to allocate track space-efficient volumes. This attribute applies only to the DS8000 storage systems.

**Storage Pool Physical Allocation Percentage**

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

### **Storage Pool Virtual Allocation Percentage**

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

### **Storage Pool Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{virtual capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB  $\div$  100 GiB  $\times$  100).

This value is only available for pools with thin-provisioned volumes.



### **Storage Pool Allocated Space (GiB)**

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

### **Storage Pool Extent Size (MiB)**

The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### **Storage Pool Tier Capacity (GiB)**

The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the amount of space on solid-state drives and hard disk drives in a pool. The report shows the storage space on solid-state drives and on hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total storage space on all of the solid-state drives and hard disk drives in a pool.

The tier is configured to be controlled by the Easy Tier function.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### **Storage Pool Tier Available Space (GiB)**

The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the amount of unused space on solid-state drives and hard disk drives in a pool. The report shows the unused space on solid-state drives and on hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total unused space on all of the solid-state drives and hard disk drives in a pool.

The tier is configured to be controlled by the Easy Tier function.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

### **Storage Pool Tier Number of MDisks**

The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the number of managed disks on the tier that are solid-state drives or hard disk drives. The

report shows the number of managed disks that are solid-state drives and the number that are hard disk drives on different rows.

- Report does not include Storage Pool Tier Type: Shows the total number of managed disks on all of the solid-state drives and hard disk drives in a tier.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

#### **Storage Pool Unallocated Volume Space (GiB)**

The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

total volume capacity - allocated space

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

#### **Storage Pool Unallocatable Volume Space (GiB)**

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

#### **Storage Pool Number of MDisks**

The number of managed disks in a storage pool.

#### **Storage Pool Compressed Virtual Capacity (GiB)**

The total virtual capacity of all the volumes that are compressed in a pool.

#### **Storage Pool Compressed Capacity (GiB)**

The amount of storage space that is used by compressed volumes in a pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

#### **Storage Pool Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

#### **Storage Pool Maximum I/O Capability**

The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Pool Back-End Storage System Type
- Storage Pool Back-End Storage RAID Level
- Storage Pool Back-End Storage Disk Type

This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

### **Storage Pool Back-End Storage Disks**

The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

### **Component properties**

You can create capacity and relationship reports that include the following information:

#### **Storage Pool RAID Level**

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

#### **Storage Pool Is Solid State**

Shows whether there are solid-state drives in the pool. This property can contain the following values:

**Mixed** The pool contains both hard disk drives and solid-state drives.

#### **Non solid state**

The pool contains no solid-state drives.

#### **Solid state**

The pool contains a least one solid-state drive.

#### **Storage Pool Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

#### **Storage Pool Format**

Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

#### **Storage Pool Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

#### **Storage Pool Logical Subsystem**

The logical subsystem (LSS) to which a volume or pool belongs.

#### **Storage Pool Rank Group**

The rank group to which a pool is assigned, that is, 0 or 1. This property applies only to the DS8000 and DS6000 storage systems.

#### **Storage Pool Lock Behavior**

Shows whether the pool is locked for write operations, or is disabled for both read and write operations when storage space is limited. This property applies only to an XIV system.

#### **Storage Pool Tier**

The tier level of pools on storage systems. For example, the tier can be 1, 2, 3, or another tier number.

#### **Storage Pool Custom Tag 1, 2, and 3**

User-defined text that is associated with a storage pool. You can add or edit the custom tags for a storage pool on the **Properties** notebook of the pool in the Tivoli Storage Productivity Center web-based GUI.

#### **Storage Pool Is Compression Active**

Shows whether the compression feature is enabled on the storage pool. If this value is Yes, the compression feature is enabled.

**Storage Pool Back-End Storage System Type**

The type of storage system that provides storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Pool Back-End Storage Disk Type**

The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Storage Pool Back-End Storage RAID Level**

The RAID level of the volumes on the back-end storage system that provide storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage Pool Status**

The condition of the resource, for example normal, warning, or error.

**Storage Pool Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Pool Tier Type**

The type of tier that is used on the storage pool. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD). The tier is configured to be controlled by the Easy Tier function.

This value is available only for SAN Volume Controller and Storwize V7000 storage systems.

**Data for storage system volumes in capacity and relationship reports**

You can include general information, capacity data, properties, information about copies of volumes, and other information about storage system volumes in capacity and relationship reports.

**Information about storage system volumes**

You can create capacity and relationship reports that include the following information:

**Storage Volume Name**

The name that was assigned to the storage volume when it was added to the system.

**Storage Volume Assigned Host Connection**

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

**Storage Volume Host Connection OS Type**

The operating system type of the server or hypervisor that the volume is assigned to.

**Storage System Volume Group ID**

The unique identifier for a collection of volumes on DS8000 and DS6000 storage systems. The storage system generates this identifier.

**Storage System Volume Group Name**

The user-defined name for a collection of volumes on DS8000 and DS6000 storage systems.

**Storage System Volume Preferred Node Name**

The name of the node that is used for I/O operations for the storage system volume.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Storage Volume Capacity (GiB)**

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

**Storage Volume Overhead (GiB)**

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

**Storage Volume Allocated Space (GiB)**

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

**Storage Volume Used Space (GiB)**

The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

**Storage Volume Used Allocated Space Percentage**

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 volumes.

**Storage Volume Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{capacity} - \text{allocated space}$$

Available only for thin provisioned volumes.

**Storage Volume Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

**Storage Volume Unused Space (GiB)**

The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, Storwize V7000, and Storwize V3700 pools.

**Storage Volume Shortfall Percentage**

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{unallocatable volume Space} \div \text{volume unallocated space}) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$$\frac{100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space}}{100 \text{ GiB unallocated volume space}} \times 100$$

**Storage Volume Unallocatable Space (GiB)**

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

**Storage Volume Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

**Storage Volume Tier Capacity (GiB)**

The amount of volume tier storage of a particular type. The type is identified in the Storage Volume Tier Type property. The tier is configured to be controlled by the Easy Tier function.

**Storage Volume Tier Available Space (GiB)**

The amount of volume tier storage of a particular type that is unused. The type is identified in the Storage Volume Tier Type property. The tier is configured to be controlled by the Easy Tier function.

## **Component properties**

You can create capacity and relationship reports that include the following information:

**Storage Volume Type**

The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume WWN**

The worldwide name of the volume.

**Storage Volume Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

**Storage Volume Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Is Assigned**

Shows whether the volume is assigned to a server. If this value is Yes, the volume is assigned to a server.

**Storage Volume Is Auto Expand Enabled**

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is Yes, the volume or volume copy can allocate new extents from a pool automatically.

**Storage Volume Is Encryptable**

Shows whether the resource can be encrypted. If this value is Yes, the resource can be encrypted.

**Storage Volume Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

**Storage Volume Is Formatted**

Shows whether a storage volume is formatted. If this value is Yes, the volume is formatted. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

**Storage Volume Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Volume Metro Mirror Name**

The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

**Storage Volume Mirror Count**

The number of mirrors that keep a synchronous copy of the resource.

**Storage Volume Throttle**

The maximum number of commands that the storage volume can queue.

**Storage Volume Warning Level**

The percentage of used capacity of the resource at which a warning is generated.

**Storage Volume FlashCopy Relationship**

Shows whether a volume on a storage system is in a FlashCopy relationship. This property can contain the following values:

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Not in a FlashCopy Relationship**

The volume is not in a FlashCopy relationship.

**Storage Volume Remote Copy Relationship**

Shows whether a volume on a storage system is in a remote copy relationship. This property can contain the following values:



**Not in a Copy Relationship**

The volume is not in a remote copy relationship.

**Source Relationship**

The volume is the source of the relationship.

**Target Relationship**

The volume is the target of the relationship.

**Both Source and Target Relationship**

The volume is both the source and target of the relationship.

**Storage Volume RAID Level**

The RAID level of a storage volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

**Storage Volume Is Compressed**

Shows whether the volume is compressed. If this value is Yes, the volume is compressed.

**Storage Volume Is Volume Control Manager Database**

Shows whether a volume is used as the Volume Control Manager Database. If this value is Yes, the volume is used as the Volume Control Manager Database.

**Storage Volume Format**

Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

**Storage Volume Serial Number**

The serial number of the resource.

**Storage Volume Number**

The number of the volume that is assigned by the system.

**Storage Volume Logical Subsystem**

The logical subsystem (LSS) to which a volume or pool belongs.

**Storage Volume Is Easy Tier**

Shows how the Easy Tier function is enabled on a volume. For example, if the value of this column is auto, the Easy Tier function is enabled in automatic mode.

**Storage Volume Tier Type**

The type of tier that is used on the volume. The tier is configured to be controlled by the Easy Tier function.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage Volume Status**

The condition of the resource, for example normal, warning, or error.

**Storage Volume Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Volume Is Easy Tier**

Shows how the Easy Tier function is enabled on a volume. For example, if the value of this column is auto, the Easy Tier function is enabled in automatic mode.

**Volume copy information**

You can create capacity and relationship reports that include the following information:

**Storage Volume Copy ID**

The identifier of a volume copy.

**Storage Volume Copy Grain Size (KiB)**

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

**Storage Volume Copy Is Auto Expand Enabled**

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is Yes, the volume or volume copy can allocate new extents from a pool automatically.

**Storage Volume Copy Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Storage Volume Copy Is Primary**

Shows whether the volume copy is the primary volume copy.

**Storage Volume Copy Is Thin Provisioned**

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is Yes, the resource is thin-provisioned.

**Storage Volume Copy Total Number of Extents**

The number of extents that are used by the storage volume copy.

**Storage Volume Copy Type**

The type of copy of the storage volume. For example, the copy of the storage volume can be sequential, a mirror copy, or another type of copy.

**Storage Volume Copy Warning Level**

The percentage of used capacity of the resource at which a warning is generated.

**Storage Volume Copy Used Allocated Space Percentage**

The percentage of reserved space for a volume copy that is being used. This value is always less than or equal to 100% because a volume copy cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$$(\text{used space} \div \text{allocated space}) \times 100$$

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 pools.

**Storage Volume Copy Used Space (GiB)**

The amount of allocated space that is used by a volume copy.

**Storage Volume Copy Allocated Space (GiB)**

The amount of space that is reserved for a volume copy.

The space that is allocated for a thin-provisioned volume copy is less than its virtual capacity, which is shown in the Storage Virtualizer Capacity (GiB) property.

For volume copies that are not thin provisioned, this value is equal to the value of the Storage Volume Copy Used Space (GiB) property.

**Storage Volume Copy Unallocated Space (GiB)**

The amount of space in a pool that is not reserved for a volume copy.

Tivoli Storage Productivity Center uses the following formula to determine this value:

capacity - allocated space

**Storage Volume Copy Physical Allocation Percentage**

The percentage of physical space that is reserved for a volume copy. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$(\text{allocated space} \div \text{capacity}) \times 100$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volume copies is 50 GiB.

**Storage Volume Copy Unused Space (GiB)**

The amount of space that is allocated to a volume copy and that is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

allocated space - used space

This property is available only for SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 pools.

**Storage Volume Copy Shortfall Percentage**

The percentage of the remaining unallocated volume copy space in a pool that is not available to be allocated to a volume copy.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

The higher the percentage value is, the more critical the shortfall of space is.

Tivoli Storage Productivity Center uses the following formula to determine this value:

$(\text{unallocatable space} \div (\text{capacity} - \text{real capacity})) \times 100$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume uses 50 GiB, 100 GiB is, nevertheless, committed to that volume (150 GiB minus 50 GiB). There is a shortfall of 30 GiB when you use the following calculation:

(100 GiB remaining commitment of volume space to the volume - 70 GiB remaining pool space)

The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

(100 GiB unallocatable volume copy space - 70 GiB remaining pool space) ÷ 100 GiB unallocated volume space × 100

This property is available only for volume copies that are thin provisioned.

#### **Storage Volume Copy Unallocatable Space (GiB)**

The unallocatable storage space in GiB of a thin-provisioned volume copy. Unallocatable space cannot be supplied by the pool to which the volume copy belongs.

#### **Storage Volume Copy Is Compressed**

Shows whether the volume copy is compressed. If this value is Yes, the volume copy is compressed.

#### **Storage Volume Copy Uncompressed Capacity (GiB)**

The amount of storage space that is used if the compressed space on the volume copy is uncompressed. For example, if 100 GiB of data is compressed, the value is 100.

### **Data for storage system disks in capacity and relationship reports**

You can include properties, and other data about storage system disks in capacity and relationship reports.

#### **Information about storage system disks**

You can create capacity and relationship reports that include the following information:

##### **Storage Disk Name**

The name that was assigned to the storage disk when it was added to the system.

#### **Capacity and usage data**

You can create capacity and relationship reports that include the following information:

##### **Storage Disk Capacity (GiB)**

The amount of storage space that is on a physical disk on the storage system.

#### **Component properties**

You can create capacity and relationship reports that include the following information:

##### **Storage Disk Flags**

Shows whether a device is a spare device. This property applies only to DS8000 storage systems.

**Storage Disk Is Encryptable**

Shows whether the resource can be encrypted. If this value is Yes, the resource can be encrypted.

**Storage Disk Is Encrypted**

Shows whether the resource is encrypted. If this value is Yes, the resource is encrypted.

**Storage Disk Is Solid State**

Shows whether the disk is a solid-state drive.

**Storage Disk Speed (RPM)**

The speed of the disk.

**Storage Disk Tag**

A number that identifies the array to which a disk belongs. This property applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems.

**Storage Disk Class**

The storage technology of the disk. This property applies only to DS8000, DS5000, and DS4000 storage systems. This property can contain the following values:

**FC** The disk is a Fibre Channel disk.

**Nearline ATA**

The disk is a nearline Advanced Technology Attachment disk.

**Nearline SAS**

The disk is a nearline serial-attached SCSI disk.

**SAS** The disk is a serial-attached SCSI disk.

**SATA** The disk is a Serial Advanced Technology Attachment disk.

**SSD** The disk is a solid-state drive.

**Unavailable**

The disk is on a DS5000 or DS4000 series system, and is not a solid-state drive.

**Unsupported**

The disk is on a system that is not a DS8000, DS5000, or DS4000 series system.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage Disk Status**

The condition of the resource, for example normal, warning, or error.

**Storage Disk Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Vendor, model, and device information**

You can create capacity and relationship reports that include the following information:

**Storage Disk Firmware Version**

The version number of the firmware that is running on the disk.

**Storage Disk Model**

The model name or model number of the resource.

**Storage Disk Serial Number**

The serial number of the resource.

**Storage Disk Vendor**

The vendor who supplied the resource.

**Data for storage system groups in capacity and relationship reports**

You can include general information about, and properties of, storage system groups in capacity and relationship reports.

**Information about storage system groups**

You can create capacity and relationship reports that include the following information:

**Storage System Group Name**

The name that was assigned to the storage system group when it was added to the system.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Storage System Group Creator**

The user name of the user who created the group.

**Storage System Group Last Modified**

The date and time that the resource was last modified.

**Storage System Group Last Modified by User**

The user name of the user who last modified the resource.

**Data for managed disks on storage systems in capacity and relationship reports**

You can include general information, capacity data, properties, and other information about managed disks on storage systems in capacity and relationship reports.

**Information about managed disks on storage systems**

You can create capacity and relationship reports that include the following information:

**Storage System MDisk Name**

The name that was assigned to the managed disk on a storage system when it was added to the system.

**Capacity and usage data**

You can create capacity and relationship reports that include the following information:

**Storage System MDisk Capacity (GiB)**

The amount of storage space on the managed disk.

**Storage System MDisk Available Space (GiB)**

The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

**Storage System MDisk Strip Size (KB)**

The RAID strip size on a managed disk on a storage system.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Storage MDisk RAID Level**

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

**Storage MDisk Spare Goal**

The number of spare drives that are required to maintain redundancy. Use spare drives to protect the system against drive failures in the array on a managed disk on a storage virtualizer.

**Storage System MDisk Type**

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

**Storage System MDisk Mode**

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

**Array** Extents are provided from local disks.

**Managed**

Extents are provided from a back-end storage volume.

**Unmanaged**

The managed disk is not used in the system.

**Storage System MDisk Is Balanced**

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is Yes, the LUNs are balanced.

**Storage System MDisk Fast Write State**

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

**Storage System MDisk Write Verify**

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is Yes, all write operations are verified by a follow-up read operation.

**Status information**

You can create capacity and relationship reports that include the following information:

**Storage System MDisk Status**

The condition of the resource, for example normal, warning, or error.

**Storage MDisk Native Status**

Shows the level of access of nodes in the system to a managed disk or volume on a storage system. The level of access can be online, offline, degraded, or excluded.

**Data for storage resource groups in capacity and relationship reports**

You can include general information about storage resource groups in capacity and relationship reports.

**Information about storage resource groups**

You can create capacity and relationship reports that include the following information:

**Storage Resource Group Name**

The name that was assigned to the storage resource group when it was created.

**Storage Resource Group Description**

A user-defined description of the storage resource group. The description is created by the user when the storage resource group is created from the Tivoli Storage Productivity Center CLI or the stand-alone GUI.

**Storage Resource Group Last Modified by User**

The user name of the user who last modified the resource.

**Storage Resource Group Custom Tag 1, 2, and 3**

User-defined text that is associated with a storage resource group.

**Data for groups in capacity and relationship reports**

You can include general information about groups in capacity and relationship reports. A group is a set of logically related volumes, file systems, and shares. For example, a group that represents a business critical application might include the volumes, file systems, and shares that provide storage to the application.

**Information about groups**

You can create capacity and relationship reports that include the following information:

**Group Name**

The name that was assigned to the group by the user when the group was created.

**Group Description**

A user-defined description of the group. The description is created by the user when the group is created from the Tivoli Storage Productivity Center CLI.

**Group Last Modified by User**

The user name of the user who last modified the resource.

**Group Custom Tag 1, 2, and 3**

User-defined text that is associated with a group.



**Group Type**

User-defined text that specifies the type of group. For example, the group type might be Application, Department, Project, General, or a similar type.

**Group Subtype**

User-defined text that is associated with a group. For example, if the Group Type property is Application, the Group Subtype property might specify the type of application. If the Group Type property is DB2 Group, the Group Subtype property might be Logs, Tablespaces, or another group that is related to DB2.

**Data for switches in capacity and relationship reports**

You can include properties and other information about switches in capacity and relationship reports.

**Information about switches**

You can create capacity and relationship reports that include the following information:

**Switch Name**

The logical name of the switch where a Fiber Channel port is located, or the name that was defined when the switch was added for monitoring. If neither name is available, Tivoli Storage Productivity Center uses the WWN of the switch.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Switch WWN**

The World Wide Name (WWN) of the switch. A WWN is the unique 64-bit identifier for the switch.

**Parent Switch Name**

The name of the parent switch of a logical switch. If a user-defined name is not set for the parent switch, the WWN is shown. If the switch is not a logical switch, no value is displayed.

**Parent Switch WWN**

The worldwide name of the parent switch.

**Switch Management Telnet Address**

The Telnet address that you access to manage the switch.

**Switch Management SNMP Address**

The Simple Network Management Protocol (SNMP) address that you access to manage the switch.

**Switch Management URL Address**

The URL that you access to manage the switch.

**Switch Domain**

The domain ID of a switch. The ID is an 8-bit identifier with a range of 0-255. This column is blank for physical switches that are parents of virtual switches.

**Switch Firmware Version**

The firmware version of the microcode on a switch.

**Switch Location**

The physical location of the switch. The location is defined when a switch is added to Tivoli Storage Productivity Center. In the web-based GUI, you can add or edit the location of the switch on the **General** tab of the properties notebook.

**Switch Custom Tag 1, 2, and 3**

User-defined text that is associated with the switch. In the web-based GUI, you can add or edit the custom tags for the switch on the **General** tab of the properties notebook.

**Switch Data Source Count**

The number of data sources that are managing a switch. A switch can be managed by multiple data sources, and a single data source can manage multiple switches. Data sources might include CIM agents, SNMP agents, or Storage Resource agents. Tivoli Storage Productivity Center communicates with a data source to collect information about a switch.

Only the data sources that were added to Tivoli Storage Productivity Center are included in this number. For example, a switch is managed by two CIM agents and one Storage Resource agent. If one of the CIM agents was not added to Tivoli Storage Productivity Center, then the value for this property is 2.

**Switch Is Virtual**

Shows whether a switch is a logical switch. If this value is Yes, the switch is a logical switch.

**Switch Ports**

The total number of ports on the switch that are not online, and ports that are online and connected to other ports.

**Switch Connected Ports**

The number of ports that are connected to a storage resource, where the storage resource can be a storage system, server, or hypervisor. This value is only available when a switch is viewed as a related resource of a storage system, server, or hypervisor.

**Status information**

You can create capacity and relationship reports that include the following information:

**Switch Status**

The condition of the resource, for example normal, warning, or error.

**Switch Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Switch Last Data Collection**

The date and time when storage statistics were last collected from the resource.

**Switch Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

## Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

### Switch IP Address

The IP address of the resource.

### Switch Serial Number

The serial number of the resource.

### Switch Vendor

The vendor who supplied the resource.

### Switch Model

The model name or model number of the resource.

## Data for switch ports in capacity and relationship reports

You can include properties and other information about switch ports in capacity and relationship reports.

### Information about switch ports

You can create capacity and relationship reports that include the following information:

#### Switch Name

The logical name of the switch where a Fiber Channel port is located, or the name that was defined when the switch was added for monitoring. If neither name is available, Tivoli Storage Productivity Center uses the WWN of the switch.

## Component properties

You can create capacity and relationship reports that include the following information:

#### Switch Port Number

The port number on the resource.

#### Switch Port WWPN

The worldwide port name of the port on the switch.

#### Switch Port Type

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N\_Port, F\_Port, or another type of port.

#### Switch Port Speed (GiB/s)

The speed of a port, which is measured in GiB per second.

#### Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

#### Switch Port Domain, Port

The domain ID of a switch, followed by the port index of a switch. In switches that have blades, the port index can differ from the port number.

## Status information

You can create capacity and relationship reports that include the following information:

**Switch Port Status**

The condition of the resource, for example normal, warning, or error.

**Switch Port Is Detected**

Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

**Switch Port Enabled State**

Shows whether a port is enabled, disabled, or is enabled but offline.

**Switch Port Operational Status**

The current operational state of the switch port. For example, this value can be Error or OK.

**Data for fabrics in capacity and relationship reports**

You can include properties and other information about fabrics in capacity and relationship reports.

**Information about fabrics**

You can create capacity and relationship reports that include the following information:

**Fabric Name**

The user-defined name for the fabric, or the name that was assigned to the fabric by the vendor.

**Fabric WWN**

The worldwide name of the fabric.

**Component properties**

You can create capacity and relationship reports that include the following information:

**Fabric Type**

The vendor of the switches in a fabric, such as IBM, Brocade, Cisco, or QLogic. If there are different switch vendors in the same fabric, the fabric type is displayed as Mixed. If this information is unknown, the property does not contain a value.

**Principal Switch of Fabric**

The name of the principal switch in the fabric at the time of the last data collection. The fabric can change its principal switch dynamically when conditions require it.

**Important:** Cisco physical switches are part of a SAN and not a fabric, therefore this property does not contain a value for those switches.

**Fabric Switches**

The number of physical and virtual switches in the fabric.

**Fabric Switch Ports**

The number of ports that are on all the switches in the fabric.

**Fabric Connected Switch Ports**

The number of switch ports that are online and connected to other ports in a fabric. The value is not shown when a fabric is viewed as a related resource of a storage system, server, or hypervisor.

**Parent Fabric Name**

The user-defined name for the parent fabric, or the name that was assigned to the parent fabric by the vendor.

**Fabric Active Zone Set**

The name of the zone set that is active. A zone set is a collection of zones within a single fabric. Only one zone set can be active for a fabric.

**Fabric Location**

The physical location of the fabric. The location is defined when a fabric is added to Tivoli Storage Productivity Center. In the web-based GUI, you can add or edit the location of the fabric on the **General** tab of the properties notebook.

**Fabric Data Source Count**

The number of data sources that are managing a fabric. A fabric can be managed by multiple data sources, and a single data source can manage multiple fabrics. Data sources might include CIM agents, SNMP agents, or Storage Resource agents. Tivoli Storage Productivity Center communicates with a data source to collect information about a fabric.

Only the data sources that were added to Tivoli Storage Productivity Center are included in this number. For example, a fabric is managed by two CIM agents and one Storage Resource agent. If one of the CIM agents was not added to Tivoli Storage Productivity Center, then the value for this property is 2.

**Fabric Is Virtual**

Shows whether a fabric is a logical fabric. If this value is Yes, the fabric is a logical fabric.

**Fabric Custom Tag 1, 2, and 3**

User-defined text that is associated with a fabric. In the web-based GUI, you can add or edit the custom tags for the fabric on the **General** tab of the properties notebook.

**Status information**

You can create capacity and relationship reports that include the following information:

**Fabric Status**

The condition of the resource, for example normal, warning, or error.

**Fabric Last Data Collection**

The date and time when storage statistics were last collected from the resource.

**Fabric Last Data Collection Status**

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

---

## Troubleshooting reports

Find answers to questions about resolving issues with reports in the IBM Cognos reporting software.

**Problem**

The Cognos software is included in Tivoli Common Reporting. When you run reports, Tivoli Common Reporting times out.

## Solution

Based on your operating system, complete the following steps to change the timeout:

### Windows

1. Open a command-line interface.
2. Stop the Jazz for Service Management server:  
`c:\program files\IBM\JazzSM\profiles\bin\stopServer.bat  
server_name -username JazzSM_username -password JazzSM_password`
3. Open the configuration directory:  
`c:\program files\IBM\JazzSM\reporting\cognos\configuration`
4. Copy the cogstartup.xml file and, then change the name of the copied file to cogstartup.orig.
5. Open the cogstartup.xml file:  
`c:\program files\IBM\JazzSM\reporting\cognos\  
configuration\cogstartup.xml`
6. In the cogstartup.xml file, enter the new timeout value in seconds for the *inactivityTimeout* setting.
7. In the cogstartup.xml file, enter new timeout value in minutes for the *mdmSessionTimeout* setting.
8. Save the cogstartup.xml file.
9. Start the Jazz for Service Management server:  
`c:\program files\IBM\JazzSM\profiles\bin\startServer.bat  
server_name`

### AIX or Linux

1. Open a command-line interface.
2. Stop the Jazz for Service Management server:  
`/opt/IBM/JazzSM/profile/bin/stopServer.sh server_name  
-username JazzSM_username -password JazzSM_password`
3. Open the configuration directory:  
`/opt/IBM/JazzSM/reporting/cognos/configuration`
4. Copy the cogstartup.xml file and, then change the name of the copied file to cogstartup.orig.
5. Open the cogstartup.xml file:  
`/opt/IBM/JazzSM/reporting/cognos/configuration/cogstartup.xml`
6. In the cogstartup.xml file, enter the new timeout value in seconds for the *inactivityTimeout* setting.
7. In the cogstartup.xml file, enter new timeout value in minutes for the *mdmSessionTimeout* setting.
8. Save the cogstartup.xml file.
9. Start the Jazz for Service Management server:  
`/opt/IBM/JazzSM/profile/bin/startServer.sh server_name`


## Problem

You can set output format options for reports either when you schedule a report or run a report. If you choose PDF format for your report and the report contains many columns, all of the columns might not fit on one page. For example, if a report contains 20 columns, then the row for a specific resource might be shown on more than one page: 10 columns on the first page and 10 columns on the second page.


### Solution

To fit the maximum number of columns on one page, set the PDF options in one of the following ways:

#### When you schedule a report

1. Navigate to a report in the Cognos software.
2. Click the **Schedule** icon .
3. Select **Override the default values**.
4. Select **PDF**, and then click **Set**.
5. From the **Orientation** list, select **Landscape**.
6. From the **Paper size** list, select **11 × 17**.

#### When you run a report

1. Navigate to a report in the Cognos software.
2. Click the **Run with options** icon .
3. Click the **advanced options** link in the tip.
4. Click **Run in the background**.
5. Select **PDF**, and then click **Set**.
6. From the **Orientation** list, select **Landscape**.
7. From the **Paper size** list, select **11 × 17**.

### Problem

When you create a custom report, you cannot find online help that describes the columns that you can include in your reports.

### Solution

To view the online help, follow these steps:

1. Go to the dashboard of the Tivoli Storage Productivity Center web-based GUI.
2. Select **Cognos > View predefined reports**.
3. Click any predefined report.
4. Click the **Help** link at the upper right of the page.
5. Go to the custom reports help content.

To view the custom reports help content in the information center, go to the product documentation at [Creating custom performance reports](#) and at [Custom reports about capacity and relationships](#).

### Problem

If the Tivoli Storage Productivity Center database (TPCDB) is not accessible, you are prompted to reenter your user ID and password although you previously entered the correct user ID and password.

### Solution

Check that the server that hosts TPCDB is available. Check that the firewall settings on the server are not preventing you from accessing TPCDB.

### Problem

If the password to access the Tivoli Storage Productivity Center database (TPCDB) is changed, you must also change the password in the Cognos software.

### Solution

In the Cognos software, complete the following steps:

1. In the **Public Folders** tab, click **Launch > Administration**.
2. Click the **Configuration** tab.
3. Click **TPCDB**, and then click **TPCDB** again.
4. In the Actions column, click the **Set properties** icon, and then click the **Signon** tab.
5. Click **Edit the signon**.
6. Make your changes, and then click **OK**.

#### Problem

When you use particular combinations of properties from the Storage Volume section in the Performance package in a custom report, a No data available message might be displayed. This occurs when you use the following combinations of properties in a custom report:

- Server Name, Hypervisor Name, and any other property from the Storage Volume section
- Server Name, Hypervisor Name, Hypervisor Cluster Name, and any other property from the Storage Volume section
- Server Name, Hypervisor Cluster Name, and any other property from the Storage Volume section

#### Solution

Do not use these combinations of properties in custom reports.

#### Problem

When you try to log in to the Cognos software, the following error message is displayed:

DPR-ERR-2109 The dispatcher cannot service the request at this time. The dispatcher is still initializing. Contact your administrator if this problem persists.

The Cognos software is included in Tivoli Common Reporting. The Jazz for Service Management server starts Tivoli Common Reporting. This problem occurs if Tivoli Common Reporting starts when DB2 is not running.

#### Solution

Stop the Jazz for Service Management server, restart DB2, and restart the Jazz for Service Management server. To do this, complete the following steps:

1. Open a command-line interface.
2. Stop the Jazz for Service Management server:

##### Windows

```
JazzSM_install_directory\profile\bin\stopServer.bat server_name
-username JazzSM_username -password JazzSM_user_password
```

##### AIX or Linux

```
JazzSM_install_directory/profile/bin/stopServer.sh server_name
-username JazzSM_username -password JazzSM_user_password
```

3. Restart DB2.
4. Start the Jazz for Service Management server:

##### Windows

```
JazzSM_install_directory\profile\bin\startServer.bat server_name
```

##### AIX or Linux


```
JazzSM_install_directory/profile/bin/startServer.sh server_name
```



**Problem**

In Internet Explorer 8, 9, or 10 when you open a report in Report Studio for the first time, a CM-REQ-4158 error might be displayed.

**Solution**

For information about how to resolve this problem, see  Report Studio does not work with Internet Explorer 8.

**Problem**

Cognos 10.2 only supports Internet Explorer 11 in compatibility mode. If you try to run Cognos reports by using Internet Explorer without setting the browser to compatibility mode, you will see error CM-REQ-4069.

**Solution**

Set the Internet Explorer 11 browser to compatibility mode.

**Problem**

When you add a filter to a property in Query Studio, not all of the resources are listed in the **Filter (Pick values from a list)** area. For example, when you add a filter to the Storage Volume Name property, not all of the volumes are listed.

This problem occurs because the maximum number of items in a list in Query Studio is 5000. If there are more than 5000 items in the list, only the first 5000 values are displayed, in alphabetical order.

**Solution**

Click **Search for values** in the **Filter (Pick values from a list)** area. Type a keyword, and then click **Search** to search for volumes on a resource. For example, type a% to find volume names that begin with 'a' or 'A'.

**Problem**

When you run a report, it takes a long time for the report output to display. It might take a long time to run a report for some combinations of the interval and reporting period. For example, the report might take a long time if you select the Sample interval or Hourly interval with one of the following periods:

- Last 30 days
- Last 90 days
- All

**Solution**

When you run a report, use shorter reporting periods with the Sample interval or Hourly interval.

**Problem**

When you run a report, the output that is displayed contains no data. If you select the Daily interval and a recent reporting period, then no data might be available. For example, if you select Last 24 hours or Today, the data for those reporting periods is unlikely to be collected yet.

**Solution**

When you run a report with a Daily interval, use a reporting period other than Last 24 hours or Today.

**Problem**

If you log off from a remote connection, you must manually restart the Jazz for Service Management server and Tivoli Common Reporting.

**Solution**

To restart the Jazz for Service Management server and Tivoli Common Reporting, complete the following steps for your operating system:

#### Windows

```
JazzSM_install_directory\profile\bin\startServer.bat server_name
```

#### AIX or Linux


```
JazzSM_install_directory/profile/bin/startServer.sh server_name
```

#### Problem

You set the default language of your operating system or your browser to Portuguese (Brazilian). However, you see English text in the report prompt pages or in the report output. This problem is a known problem with IBM Cognos 10.2 and Tivoli Common Reporting Version 3.1.0.2.

#### Solution

To set the language of the report prompt pages and report output to Portuguese (Brazilian), complete the following steps:

1. In the Cognos software, click the **My Area Options** icon  .
2. Click **My Preferences**.
3. In the **Regional options** area, under **Content language**, select **Use the following language**.
4. Select **Portuguese (Brazil)**.
5. Click **OK**.

#### Problem

If you create a custom performance report in Query Studio, in particular circumstances, the report takes a long time to complete. This problem might occur in the following circumstances:

- The report is for more than 10,000 volumes.
- The interval for the report is Hourly, and the data for the report was collected for more than 15 days.
- The interval for the report is Sample, and the data for the report was collected for more than 7 days.

#### Solution

When this problem occurs, complete the following steps for the report in Query Studio:

1. Click **Run Report** in the menu pane.
2. Click **Advanced Options**.
3. Clear each of the following check boxes:
  - **Automatically generate footer summaries for measures**
  - **Automatically summarize detail values, suppressing duplicates**
  - **Enable drill up and drill down in the report output**
  - **Enable drill through from a package in the report output**
4. Click **OK**.

#### Problem

Some characters might not be displayed correctly in reports in PDF format. Characters from the Mongolian, Tibetan, or Yi languages are displayed as square boxes in reports in PDF format.

#### Solution

To resolve this problem, you must enable IBM Cognos to display Chinese characters correctly. For more information about how to enable IBM

Cognos to display Chinese characters correctly, go to the product documentation at [Enabling the Proper Display of Chinese Characters](#).

---

## Showing the package version number and build ID

Access information about the packages for custom reports and predefined reports, such as the version number and the build ID.

You can identify the version number and build ID for each of the following packages:

- Capacity and Relationships
  - Performance
1. Start the web-based GUI for Tivoli Storage Productivity Center.
  2. Depending on the package whose version number and build ID you want to see, complete one of the following steps:
    - To identify the version number and build ID of the Capacity and Relationships package, select **Cognos > Create capacity reports**.
    - To identify the version number and build ID of the Performance package, select **Cognos > Create performance reports**.
  3. In the navigation tree, position the cursor over the item under the name of the package.

The version number and the build ID for the package is shown.

---

## Reporting in a vSphere environment

You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices.

You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems.

To view Tivoli Storage Productivity Center storage information in block and file storage reports, you must register Tivoli Storage Productivity Center as a VASA provider.

For VASA provider-specific reports, ensure that the storage systems from which hypervisors have volumes or assigned shares are managed by Tivoli Storage Productivity Center.

For **Fabric Connections**, ensure that both the hypervisor and the fabric that is part of are managed by Tivoli Storage Productivity Center.

For **Storage Mappings**, ensure that both the hypervisor and the storage system are managed by Tivoli Storage Productivity Center and that there is an SRA deployed on the virtual machine.

For **Storage System Metrics**, ensure that both the hypervisor and the storage system are managed by Tivoli Storage Productivity Center and that performance is running on the storage system.

Use the following vSphere reports to view information about your virtual resources and the back-end storage systems:

### Fabric Connections

This storage report displays fabric information that includes zone and switch details in the vSphere Web Client.

### Storage Mapping

This storage mapping report displays end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and the virtual resources that are monitored by vSphere.

### Storage System Metrics

This performance report displays performance metrics that include the total I/O rate and response time for the back-end storage systems that are monitored by Tivoli Storage Productivity Center, and that performance is running on the storage system.

### SCSI Volumes (LUNs)

This volume report displays block storage information that is provided by Tivoli Storage Productivity Center. The information includes the following details:

- Space that is committed to the volume
- Thin provisioning status
- System capability
- Storage array name
- Volume identifier on the array
- Namespace of the VASA provider

### Datastores

The Datastores report includes file system information that is provided by Tivoli Storage Productivity Center, including the system capability and the namespace of the VASA provider.

### NAS Mounts


The NAS Mounts report includes file system information that is provided by Tivoli Storage Productivity Center, including the system capability and the namespace of the VASA provider.




## Viewing vSphere storage mapping to Tivoli Storage Productivity Center

Use the vSphere Web Client extension to view the end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and virtual resources that are monitored by vSphere.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select the appropriate inventory object. For example, select a virtual machine.
2. Click the **Monitor** tab to view the monitoring options for the object.
3. Click the **Storage Mapping** tab to view a table that contains the mapping information from the virtual resource to Tivoli Storage Productivity Center.
4. Optional: To change the column that determines the sort order of resources, click the name of that column. For example, click the **Virtual Disk Size** column to sort the list of virtual disks according to size.

**Tip:** To view descriptions of the columns in the **Storage Mapping** table, click the **Help** icon  beside the **Hypervisor Host Name**.

5. Optional: To retrieve new data from Tivoli Storage Productivity Center, click the **Refresh** icon  beside the **Guest OS Host Name**. The **Last updated** information in the toolbar is also updated.
6. Optional: To export storage mapping data to a CSV file or to copy the data to the clipboard, click the **Export** icon  or select the relevant option from the **Export menu**  in the status bar.

### vSphere storage mapping to Tivoli Storage Productivity Center

You can use the vSphere Web Client extension to view the end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and virtual resources that are monitored by vSphere.

In the vSphere Web Client, use the **Monitor** tab to access the **Storage Mapping** tab for a virtual resource. You can view the attributes of each virtual disk, including the information that is provided by Tivoli Storage Productivity Center.

Use the **Storage Mapping** tab to help you to understand the relationships between an inventory object and all the virtual and physical storage resources that are related to this object. You can view details of storage usage inside the Guest Operating System (OS) and the back-end storage details such as volume, pool, and storage system information.

The Storage attribute column in the following table displays all of the attributes that are displayed in the **Storage Mapping** tab for a virtual disk. The Information source column shows whether the attribute information is from Tivoli Storage Productivity Center or vSphere.

*Table 110. vSphere Storage mapping to Tivoli Storage Productivity Center*

Storage attribute	Information source
Guest OS Host Name	vSphere
Virtual Disk	vSphere
Virtual Disk Size	vSphere
Virtual Disk Maximum Size	vSphere
Thin Provisioned	vSphere
Raw Device Mapping	vSphere
File	vSphere
Mount Point on Guest OS	Tivoli Storage Productivity Center  This value is not available and "--" is displayed in any of the following situations: <ul style="list-style-type: none"> <li>• The Guest OS on the virtual machine is not monitored by Tivoli Storage Productivity Center using a Storage Resource agent.</li> <li>• The ESX Server is not monitored by Tivoli Storage Productivity Center.</li> <li>• There is no file system on the virtual disk.</li> </ul>

Table 110. vSphere Storage mapping to Tivoli Storage Productivity Center (continued)

Storage attribute	Information source
Used Space on Guest OS	<p>Tivoli Storage Productivity Center</p> <p>This value is not available and "--" is displayed in any of the following situations:</p> <ul style="list-style-type: none"> <li>• The Guest OS on the virtual machine is not monitored by Tivoli Storage Productivity Center using a Storage Resource agent.</li> <li>• The ESX Server is not monitored by Tivoli Storage Productivity Center.</li> <li>• There is no file system on the virtual disk.</li> </ul>
Disk on Guest OS	<p>Tivoli Storage Productivity Center</p> <p>This value is not available and "--" is displayed in any of the following situations:</p> <ul style="list-style-type: none"> <li>• The Guest OS on the virtual machine is not monitored by Tivoli Storage Productivity Center using a Storage Resource agent.</li> <li>• The ESX Server is not monitored by Tivoli Storage Productivity Center.</li> <li>• There is no file system on the virtual disk.</li> </ul>
Disk on Hypervisor	vSphere
Storage System	<p>Tivoli Storage Productivity Center</p> <p>This value is not available and "--" is displayed if the storage system or ESX Server is not monitored by Tivoli Storage Productivity Center.</p> <p>This value is not applicable and "N/A" is displayed if the back-end storage is a NAS device.</p>
Storage Pool	<p>Tivoli Storage Productivity Center</p> <p>This value is not available and "--" is displayed if the storage system or ESX Server is not monitored by Tivoli Storage Productivity Center.</p> <p>This value is not applicable and "N/A" is displayed if the back-end storage is a NAS device.</p>
Storage Volume	<p>Tivoli Storage Productivity Center</p> <p>This value is not available and "--" is displayed if the storage system or ESX Server is not monitored by Tivoli Storage Productivity Center.</p> <p>This value is not applicable and "N/A" is displayed if the back-end storage is a NAS device.</p>

Table 110. vSphere Storage mapping to Tivoli Storage Productivity Center (continued)

Storage attribute	Information source
NAS Server	<p>vSphere</p> <p>This value is not available and "--" is displayed if the NAS server or ESX Server is not monitored by Tivoli Storage Productivity Center.</p> <p>This value is not applicable and "N/A" is displayed if the back-end storage is a storage system.</p>
NAS Share Path	<p>vSphere</p> <p>This value is not available and "--" is displayed if the NAS server or ESX Server is not monitored by Tivoli Storage Productivity Center.</p> <p>This value is not applicable and "N/A" is displayed if the back-end storage is a storage system.</p>

#### Related tasks:

“Viewing vSphere storage mapping to Tivoli Storage Productivity Center” on page 1014

Use the vSphere Web Client extension to view the end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and virtual resources that are monitored by vSphere.

#### CSV file exports

You can use the export actions to export the information from the reports tables in the vSphere Web Client extension for Tivoli Storage Productivity Center to a CSV file. You can also copy the information to the clipboard.

Columns that are hidden on a page are included in the generated file. Depending on your web browser, you can determine the name and location of the file. The default file name is in the format `export_yyyy_mm_dd_hh_mm_ss.csv`.

To export information, click the **Export** icon  or select the relevant option from the **Export menu**  in the status bar. If no resources are selected, the information for all resources is exported.

## Viewing storage system metrics in vSphere Web Client extension

Use the vSphere Web Client extension to view the key performance data for virtual resources that are monitored by vSphere. The data is retrieved from the back-end storage systems that are monitored by IBM Tivoli Storage Productivity Center.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters** and browse to the host for which you want to view performance data.

2. Click **Monitor > Performance > Storage System Metrics** to navigate to the performance data that was collected by Tivoli Storage Productivity Center. The performance information is displayed in the following formats:


#### **Performance Chart**

The chart shows a visual representation of how storage resource performance trends over time. Each line on the chart represents a metric and a resource. If the corresponding volumes have the same metric value, lines are overlaid on the chart.




#### **Performance Chart Legend**

The Performance Chart Legend shows performance information that is formatted into rows and columns.

3. To show the performance chart for a volume, in the **Performance Chart Legend**, click the volume. To select multiple volumes, press Ctrl and click the volumes. To select a series of volumes, select the first volume, and then press Shift and click the last volume. You can also press Shift and the Up/Down Arrow. If you select more than 10 volumes, only the first 10 volumes are shown in the performance chart. When you select volumes, the performance chart for those volumes is shown.
4. To select the time period for which you want to view performance data, select a value in the **Time** field. New data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.
5. To select the metric that you want to view, select a value in the **Metric** field. New data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.
6. To view the default chart for the 10 volumes with the worst performance, press Ctrl and click the selected volumes to deselect them. The chart is generated again.

**Tip:** To view descriptions of the metrics in the performance chart and the columns in the **Performance Chart Legend**, click the **Help** icon  beside the

**Refresh** icon .

7. Optional: To retrieve new data from Tivoli Storage Productivity Center, click the **Refresh** icon . The **Last updated** information in the toolbar is also updated.
8. Optional: To customize the column display in the **Performance Chart Legend**, use the pop-up menu for the table.
9. Optional: To export storage system metrics to a CSV file or to copy the data to the clipboard, click the **Export** icon  or select the relevant option from the **Export menu**  in the status bar.

#### **Related reference:**

“Report Displays” on page 1024

You can change the appearance of the **Storage System Metrics** table in the vSphere Web Client extension for Tivoli Storage Productivity Center. You can customize the appearance of table columns, including showing columns that are hidden, and hiding columns that are shown.

“CSV file exports” on page 1017

You can use the export actions to export the information from the reports tables in the vSphere Web Client extension for Tivoli Storage Productivity Center to a CSV file. You can also copy the information to the clipboard.



## Storage System Metrics

Use the **Storage System Metrics** tab to view key performance data about the storage system that was collected from IBM Tivoli Storage Productivity Center during the last hour, day, or week. You can view the data in chart or table format, and export data to a CSV file. If no performance data was collected for the resource, the message No data available is displayed.

### Performance Chart


You can view the key performance data for storage systems in a chart layout. The chart shows a visual representation of how storage resource performance trends over time. The x-axis shows the time period in which the data was collected. The y-axis shows the unit of measurement for a metric. Each line on the chart represents a metric and a resource. By default, the chart displays lines that represent the read I/O rate for the 10 volumes with the worst performance in the last hour.

**Tip:** Lines are overlaid on the chart if the corresponding volumes have the same metric value.

Click a row in the **Performance Chart Legend** to show the performance chart for that volume. To select multiple volumes, press Ctrl and click the volumes. To select a series of volumes, select the first volume, and then press Shift and click the last volume. You can also press Shift and the Up/Down Arrow. To deselect a volume, press Ctrl and click the selected volume.

**Restriction:** If you select more than 10 volumes, only the first 10 volumes are shown in the performance chart.

When you deselect all volumes, the default chart for the 10 volumes with the worst performance is generated again.

**Note:** When you select or deselect a volume, the chart is updated, but new data is not retrieved from Tivoli Storage Productivity Center. To retrieve new data, click the **Refresh** icon .

You can view the following information for volumes in the performance chart:

**Time** The time period that is used to select performance data. When you change the time period, new data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.

You can select from the following options:

#### Last 1 Hour

The performance data that is selected on a sample basis from the data that was collected over the last hour. This sample data is collected at the same frequency as the data collection by the performance monitor for the volume. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments for the last hour. The default time period is **1 hour**.

#### Last 1 Day

The performance data that is selected on an hourly basis from the data that was collected over the last day.

### **Last 1 Week**

The performance data that is selected on an hourly basis from the data that was collected over the last week.

**Metric** The performance metric that is displayed for the selected volume or volumes. When you change the **Metric**, new data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.

You can select from the following options:

#### **Read I/O rate (ops/s)**

The average number of read operations per second over the specified time period.

#### **Write I/O rate (ops/s)**

The average number of write operations per second over the specified time period.

#### **Total I/O rate (ops/s)**

The average number of read and write operations per second over the specified time period.

#### **Read data rate (KiB/s)**

The average number of kibibytes per second that were transferred for read operations over the specified time period.

#### **Write data rate (KiB/s)**

The average number of kibibytes per second that were transferred for write operations over the specified time period.

#### **Total data rate (KiB/s)**

The average number of kibibytes per second that were transferred for read and write operations over the specified time period.

#### **Read response time (ms/op)**

The average number of milliseconds to complete a read operation over the specified time period.

#### **Write response time (ms/op)**

The average number of milliseconds to complete a write operation over the specified time period.

#### **Overall response time (ms/op)**

The average number of milliseconds to complete read and write operations over the specified time period.

#### **Read transfer size (KiB/op)**

The average number of kibibytes that are transferred per read operation over the specified time period.


#### **Write transfer size (KiB/op)**

The average number of kibibytes that are transferred per write operation over the specified time period.

#### **Overall transfer size (KiB/op)**

The average number of kibibytes that are transferred for read operations and write operations over the specified time period.

### **Refresh**

Click the **Refresh** icon  to retrieve new data from Tivoli Storage Productivity Center and generate a new chart.

Hover the mouse pointer over a point on a volume line to view the following information:

**Volume WWN**

The worldwide name of the volume.

**Time** The date and time that the performance data was collected.

The details for the selected metric are also displayed. If you hover the mouse pointer over the highest point on a volume line, the highest (worst) value for the **I/O rate**, **Data rate**, or **Response time** metric, is displayed. If the **Transfer size** metric is selected, the worst value is the lowest point on the volume line.

## Performance Chart Legend

You can view performance data for storage systems in table format. The table shows performance information that is formatted into rows and columns. Each row represents a resource and a time stamp. Each metric is shown as a column in the table. The type of resource determines the metrics and information that is available. Click a column name to change the column that is used to sort the table. Click the link to open the associated page in the Tivoli Storage Productivity Center web-based GUI.

You can view information for each volume that is monitored by Tivoli Storage Productivity Center. The following information is displayed in the table by default:

**Key** The color key for the volume. The chart line for a volume is displayed in this color.

**Tip:** The chart color palette uses 24 colors. If the number of volumes is greater than 24, the colors repeat.

**Name** The name of the volume. Click the link to open the **Hypervisors > View Details > Disks > View Properties** page in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

**Tip:** Right-click on any column heading in the volumes table and select the **Unique ID** column. Use the information in this column to filter the search to locate the volume in the table.

**Location**

The location of the volume.

**Snapshot time**

The date and time that the performance data for the volume was collected.

**Interval (sec.)**

The size of the sample interval (in seconds) for the collection of performance data for the volume.

**Storage system**

The name of the storage system as displayed in Tivoli Storage Productivity Center. Click the link to open the **Storage System > Block Storage > View Details** page for the storage system in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

**Storage volume**

The name or identifier of the storage volume as displayed in Tivoli Storage

Productivity Center. Click the link to open the **Storage System > Block Storage > View Details > Volume Properties** page for the storage volume in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

**Storage pool**

The name of the storage pool as displayed in Tivoli Storage Productivity Center. Click the link to open the **Storage System > Block Storage > View Details > Pool Properties** page for the storage pool in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

**Read I/O rate sample (ops/s)**

The average number of read operations per second for the last hour of sample performance data.

**Read I/O rate hourly (ops/s)**

The hourly average number of read operations per second for the selected last day or week of performance data.

**Total I/O rate sample (ops/s)**

The average number of read and write operations per second for the last hour of sample performance data.

**Total I/O rate hourly (ops/s)**

The hourly average number of read and write operations per second for the selected last day or week of performance data.

**Overall response time sample (ms/op)**

The average number of milliseconds that it took to complete each read and write operation for the latest sample of performance data.

**Overall response time hourly (ms/op)**

The hourly average number of milliseconds that it took to complete each read and write operation for the selected last day or week of performance data.

When you select a **Metric** in the performance chart, the corresponding columns for both sample and hourly averages are also displayed in the table:

**Write I/O rate sample (ops/s)**

The average number of write operations per second for the last hour of sample performance data.

**Write I/O rate hourly (ops/s)**

The hourly average number of write operations per second for the selected last day or week of performance data.

**Read data rate sample (KiB/s)**

The average number of kibibytes per second that were transferred for read operations for the latest sample of performance data.

**Read data rate hourly (KiB/s)**

The hourly average number of kibibytes per second that were transferred for read operations for the selected last day or week of performance data.

**Write data rate sample (KiB/s)**

The average number of kibibytes per second that were transferred for write operations for the latest sample of performance data.

**Write data rate hourly (KiB/s)**

The hourly average number of kibibytes per second that were transferred for write operations for the selected last day or week of performance data.

**Total data rate sample (KiB/s)**

The average number of kibibytes per second that were transferred for read and write operations for the latest sample of performance data.

**Total data rate hourly (KiB/s)**

The hourly average number of kibibytes per second that were transferred for read and write operations for the selected last day or week of performance data.

**Read response time sample (ms/op)**

The average number of milliseconds that it took to service each read operation for the latest sample of performance data.

**Read response time hourly (ms/op)**

The hourly average number of milliseconds that it took to service each read operation for the selected last day or week of performance data.

**Write response time sample (ms/op)**

The average number of milliseconds that it took to service each write operation for the latest sample of performance data.

**Write response time hourly (ms/op)**

The hourly average number of milliseconds that it took to service each write operation for the selected last day or week of performance data.

**Read transfer size sample (KiB/op)**

The average number of kibibytes that were transferred per read operation for the latest sample of performance data.

**Read transfer size hourly (KiB/op)**

The hourly average number of kibibytes that were transferred per read operation for the selected last day or week of performance data.

**Write transfer size sample (KiB/op)**

The average number of kibibytes that are transferred per write operation for the latest sample of performance data.

**Write transfer size hourly (KiB/op)**

The hourly average number of kibibytes that are transferred per write operation for the selected last day or week of performance data.

**Overall transfer size sample (KiB/op)**

The average number of kibibytes that were transferred per read and write operation for the latest sample of performance data.

**Overall transfer size hourly (KiB/op)**

The hourly average number of kibibytes that were transferred per read and write operation for the selected last day or week of performance data.

**Related reference:**

“Report Displays” on page 1024

You can change the appearance of the **Storage System Metrics** table in the vSphere Web Client extension for Tivoli Storage Productivity Center. You can customize the appearance of table columns, including showing columns that are hidden, and hiding columns that are shown.

“CSV file exports” on page 1017

You can use the export actions to export the information from the reports tables in the vSphere Web Client extension for Tivoli Storage Productivity Center to a CSV

file. You can also copy the information to the clipboard.

### Report Displays:

You can change the appearance of the **Storage System Metrics** table in the vSphere Web Client extension for Tivoli Storage Productivity Center. You can customize the appearance of table columns, including showing columns that are hidden, and hiding columns that are shown.

Use the pop-up menu for the table to show, hide, size, or lock columns in the table:

#### Hide [Column Name]

Hides the selected column in the table.

**Restriction:** If only one column is visible, this column cannot be hidden.

#### Size [Column Name] to Fit

Resizes the selected column to fit the data in the column.

#### Size All Columns to Fit

Resizes all columns in the table to fit in the view and removes the scrollbar.

#### Lock First Column

Locks the first available column in the table and keeps it visible when you are scrolling.

#### Show/Hide Columns

Shows or hides the columns in the table.

#### Show Toolbar


Changes the display space for the table by showing or hiding the performance chart toolbar. If you hide the toolbar, you cannot change the metric or time that is displayed in the chart. You also cannot refresh the data and generate a new chart. Show the toolbar again to use these options.

## Viewing Tivoli Storage Productivity Center fabric connections in vSphere Web Client extension

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to view information about the fabric that is connected to a storage adapter.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select a host and click the **Manage** tab.
3. Click the **Storage** tab, and then click **Storage Adapters**.
4. In the **Adapter Details** section of the **Storage Adapters** page, select a Fibre Channel host bus adaptor.
5. Click the **Fabric Connections** tab to view the details of the fabric that the host is connected to.

**Tip:** To view descriptions of the fields in the **Fabric Connections** tab, click the **Help** icon  in the upper-right part of the tab.

6. Optional: Click the links in the Fabric name, Switch name, and Switch port fields to open the associated pages in the Tivoli Storage Productivity Center

web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the pages are displayed.

## Viewing block storage reports in the vSphere Web Client

Use the vSphere Web Client to view information for volumes. When you configure Tivoli Storage Productivity Center as a VASA provider, you can view detailed block storage information in the SCSI Volumes (LUNs) report. When you use the Provision Block Storage wizard to create a data store, you can view the details in the Datastores report.

**Tip:** You can also view detailed block storage information from Tivoli Storage Productivity Center in the vSphere Client.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**, and browse to the entity that contains the volumes or data stores that you want to view. You can select from the following entities:
  - Data center
  - Host
  - Cluster
  - Virtual machine

The SCSI Volumes (LUNs) report is not available for data store clusters.

2. Click **Monitor > Storage Reports**.
3. To view the details for all volumes, in the **Report On** list, select the **SCSI Volumes (LUNs)** category.
4. To view the details for all data stores, in the **Report On** list, select the **Datastores** category.
5. Optional: Move the cursor over the column headings to see the description of each column. An asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center.

### Block storage reports in vSphere Web Client

You can view block storage reports in the vSphere Web Client. When you configure Tivoli Storage Productivity Center as a VASA provider, you can then view detailed block storage information in the storage reports. When you use the Provision Block Storage wizard to create a data store, you can view the details in the Datastores report.

Use the **Monitor > Storage Reports** page to view the reports. In the vSphere Web Client, an asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center.

### SCSI Volumes (LUNs) report

You can view the following information, which is provided by Tivoli Storage Productivity Center, in the SCSI Volumes (LUNs) report:

#### Committed

The amount of storage space that is allocated to the volume by the back-end storage system, if the volume is thin-provisioned.

#### Thin Provisioned

Shows whether the volume is thin provisioned.

**System Capability**

The storage capability of the volume. This capability describes the physical capabilities and services that the back-end storage offers, such as performance and redundancy of the volume.

**Storage Array**

The identifier of the storage array.

**Identifier on Array**

The identifier of the volume on the storage array.

**Storage Provider Namespace**

The namespace of the storage provider.

**Datastores report**

You can view the following information, which is provided by Tivoli Storage Productivity Center, in the Datastores report:

**System Capability**

The storage capability of the data store. This capability describes the physical capabilities and services that the back-end storage offers, such as performance and redundancy of the data store.

**User Defined Capability**

The namespace of the storage provider.

**Storage Provider Namespace**

The namespace of the storage provider.

## Viewing file system reports in the vSphere Web Client

Use the vSphere Web Client to view detailed information for file systems. When you configure Tivoli Storage Productivity Center as a VASA provider, you can then view detailed file system information in the vSphere Web Client Datastores report and in the NAS Mounts report.

**Tip:** You can also view detailed file system information from Tivoli Storage Productivity Center in the vSphere Client.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**, and browse to the entity that contains the NAS mounts that you want to view. You can select from the following entities:
  - Data center
  - Host
  - Cluster
  - Virtual machine

The NAS Mounts report is not available for data store clusters.

2. Click **Monitor > Storage Reports**.
3. To view the details for all data stores, in the **Report On** list, select the **Datastores** category.
4. To view the details for all NAS mounts, in the **Report On** list, select the **NAS Mounts** category.
5. Optional: Move the cursor over the column headings to see the description of each column. An asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center VASA provider.



## File system reports in vSphere Web Client

You can view storage reports in the vSphere Web Client. When you configure Tivoli Storage Productivity Center as a VASA provider, you can then view detailed file system information in the storage reports.

Use the **Monitor > Storage Reports** page to view the **Datastores** report and the **NAS Mounts** report. In the vSphere Web Client, an asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center.

You can view the following information that is provided by Tivoli Storage Productivity Center, in both the **Datastores** report and the **NAS Mounts** report:

### System Capability

The storage capability of the NAS mount. For file system storage, this column displays either the service class that is associated with the NAS mount or the value File.

### Storage Provider Namespace

The namespace of the storage provider.

## Viewing Tivoli Storage Productivity Center VASA provider reports

Use the Tivoli Storage Productivity Center VASA provider to view Tivoli Storage Productivity Center data in vCenter Server reports.

When a Tivoli Storage Productivity Center VASA provider is online and inventory details are loaded, you can use the vSphere Client or the vSphere Web Client to view Tivoli Storage Productivity Center data in vCenter Server reports. vCenter Server periodically updates storage data from the Tivoli Storage Productivity Center VASA provider. You can start an update to ensure that you can view current data.

### Viewing Tivoli Storage Productivity Center VASA provider reports in the vSphere Client

Use the Tivoli Storage Productivity Center VASA provider to view Tivoli Storage Productivity Center data in the vSphere Client reports.

When a Tivoli Storage Productivity Center VASA provider is online and inventory details are loaded, you can use the vSphere Client to view Tivoli Storage Productivity Center data in vCenter Server reports. vCenter Server periodically updates storage data from the Tivoli Storage Productivity Center VASA provider. You can start an update to ensure that you can view current data.

1. Start the vSphere Client, and log on to the vCenter Server system.
2. Select a host, and click **Storage Views > Reports**.
3. To view storage data, in the **Show all** list, select the **Show all SCSI Volumes (LUNs)**, **Show all Datastores**, or **Show all NAS Mounts** category.

**Tip:** Storage views are updated automatically to show recently added or deleted volumes or shares.

4. Optional: To ensure that you are viewing current storage data for other storage views, get the most recent vCenter Server updates from the Tivoli Storage Productivity Center VASA provider by completing the following steps:
  - a. Select **View > Administration > Storage Providers**.
  - b. Select the Tivoli Storage Productivity Center VASA provider, and click **Sync**.

- c. Allow some minutes for the synchronization to take effect. Click **Refresh All**. Verify that the **Last Sync Time** value for the VASA provider is updated to the latest time stamp.
- d. From **Home**, click **Hosts and Clusters**.
- e. Select **ESX** and click the **Storage Views** tab.
- f. Click the **Update...** link to update the storage views.

### Viewing Tivoli Storage Productivity Center VASA provider reports in the vSphere Web Client

Use the Tivoli Storage Productivity Center VASA provider to view Tivoli Storage Productivity Center data in the vSphere Web Client reports.

When a Tivoli Storage Productivity Center VASA provider is online and inventory details are loaded, you can use the vSphere Web Client to view Tivoli Storage Productivity Center data in vCenter Server reports. vCenter Server periodically updates storage data from the Tivoli Storage Productivity Center VASA provider. You can start an update to ensure that you can view current data.

1. Start the vSphere Web Client and log on to the vCenter Server system.
2. From the Inventories view, select **Hosts and Clusters**.
3. Select a host, click the **Monitor** tab, and then click **Storage Reports**.
4. From the **Report On** list, select **Show all SCSI Volumes (LUNs)**, **Show all Datastores**, or **Show all NAS Mounts**.

**Tip:** Storage views are updated automatically to show recently added or deleted volumes or shares.

5. Optional: To ensure that you are viewing current storage data for other storage views, get the most recent vCenter Server updates from the Tivoli Storage Productivity Center VASA provider by completing the following steps:
  - a. From the Inventories view, select **Hosts and Clusters**.
  - b. In the **Manage** tab, click **Storage Providers**.
  - c. Select the Tivoli Storage Productivity Center VASA provider, and click the **Rescan the storage provider** icon.
  - d. Allow some minutes for the task to complete. Verify that the **Last Rescan Time** value for the VASA provider is updated to the latest time stamp.
  - e. From the Inventories view, select **Hosts and Clusters**.
  - f. Select a host, click the **Monitor** tab, and then click **Storage Reports**.
  - g. Click the **Rescan vCenter Server for new inventory and storage devices** icon to update the storage reports.

---

## Reporting with the Tivoli Storage Productivity Center Monitoring Agent

Use Tivoli Enterprise Portal to view a set of consolidated reports for Tivoli Storage Productivity Center. Tivoli Enterprise Portal is the interface for viewing information that is collected by the Tivoli Storage Productivity Center Monitoring Agent.

**Tip:** Before you can use Tivoli Enterprise Portal to view information about Tivoli Storage Productivity Center, you must install and configure the Tivoli Storage Productivity Center Monitoring Agent. For more information, see

 [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.doc/fqz0\\_c\\_installing\\_itm.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.doc/fqz0_c_installing_itm.html)

After the Monitoring Agent is installed and configured, the following predefined workspaces are available on the Tivoli Enterprise Portal window:

### TPC Server Overview

Provides an overview of the key parameters and infrastructure metrics about the Tivoli Storage Productivity Center server. The Tivoli Storage Productivity Center server consists of the Data server and Device server.

### TPC Data Collection Overview

Displays the status of the most recent probes and performance monitors.

### Managed SAN Overview

Provides an overview of the metrics that measure the performance of monitored resources in a SAN environment.

### Managed SAN Health

Monitors the health and availability of the SAN-attached resources that are managed by Tivoli Storage Productivity Center.

## TPC Server Overview workspace

Use this workspace to view key parameters and infrastructure metrics about the Data server and Device server. To access this workspace, click **TPC Server Overview** in the navigation tree of the Navigator window.

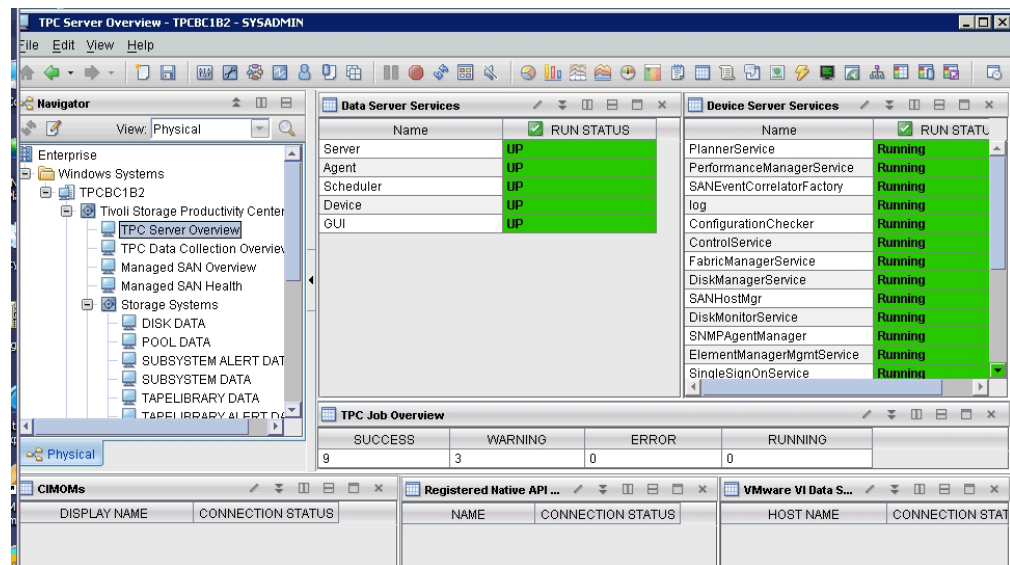


Figure 6. TPC Server Overview report

This workspace provides the following views:

### Data Server Services

Displays a table that shows all the Data server services (Server, Agent, Scheduler, Device, and GUI) including their run status.

### Device Server Services

Displays a table that shows all Device server services, including their run status.

### TPC Job Overview

Displays a table that shows statistics about the total number of jobs and their corresponding status.

## CIMOMs

Displays a table that shows all the managed CIM agents, including their connection status.

## Registered Native API Devices

Displays a table that shows all managed native interface devices, including their corresponding connection status.

## VMware VI Data Sources

Displays a report that shows all attached VMware VI Data Sources and their corresponding connection status.

## TPC Data Collection Overview workspace

Use this workspace to view the status of the latest probes and performance monitors. To access this workspace, click **TPC Data Collection Overview** in the navigation tree of the **Navigator** window.

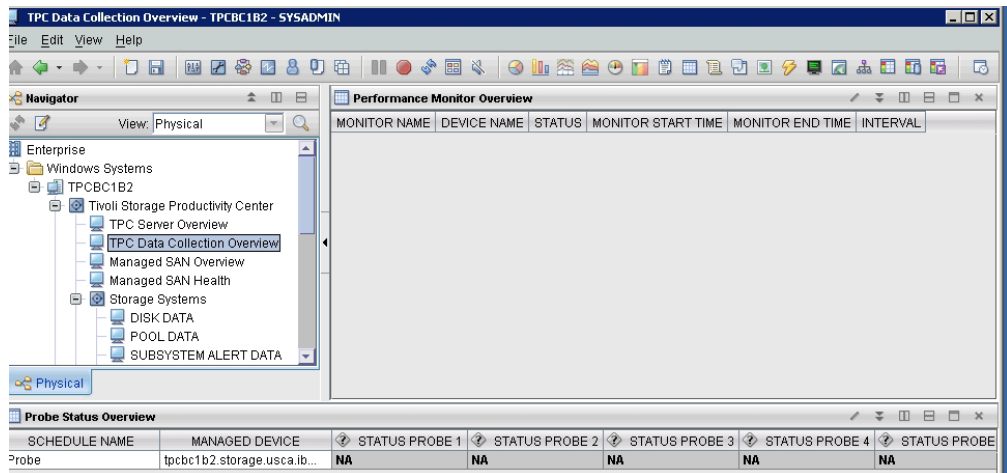


Figure 7. TPC Data Collection Overview report

This workspace provides the following information:

### Probe status overview

Displays a list probes and the status of their last five runs.

### Performance monitor overview

Displays a list of performance monitors and their corresponding status.

## Managed SAN Overview workspace

Use the Managed SAN Overview workspace to view metrics that measure the performance of monitored resources in your SAN environment. To access this workspace, click **Managed SAN Overview** in the navigation tree of the Navigator window.

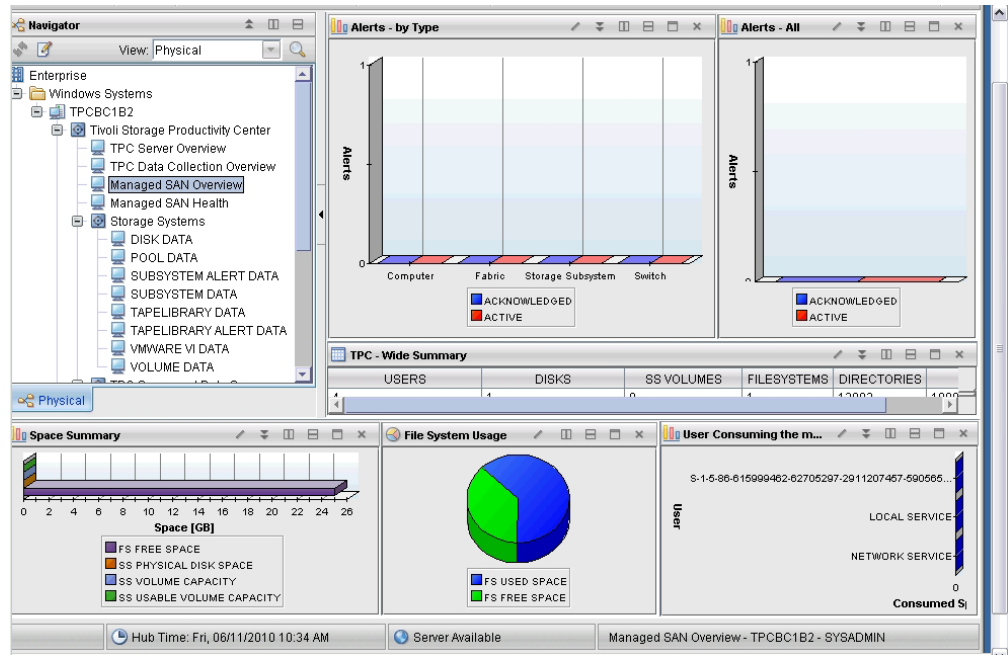


Figure 8. Managed SAN overview report

You can see the following views:

#### Alerts - by Type

Displays a bar chart that shows the number of all computers, switches, storage systems, and fabric alerts, including their current state (acknowledged or active).

#### Alerts - All

Displays a bar chart that shows the number of alerts for all resources, including their current state (acknowledged or active).

#### TPC-Wide Summary

Displays the total amount of monitored users, disks, storage system volumes, file systems, files, and folders.

#### Space Summary

Displays a bar chart that shows the total amount of file system capacity, physical disk space of all storage systems, total amount of the available volume capacity, and total amount of usable volume capacity in the managed storage environment.

#### User Consuming the most Space

Displays a chart with the total amount of space that is allocated by single users. Only the top 10 users who consume the most space is shown.

#### File System Usage

Displays a pie chart that shows the total amount of free and used file system space in the managed storage environment.

### Managed SAN Health workspace

Use this workspace to view a health summary of the resources in the managed SAN environment. To access this workspace, click **Managed SAN Health** in the navigation tree of the Navigator window.

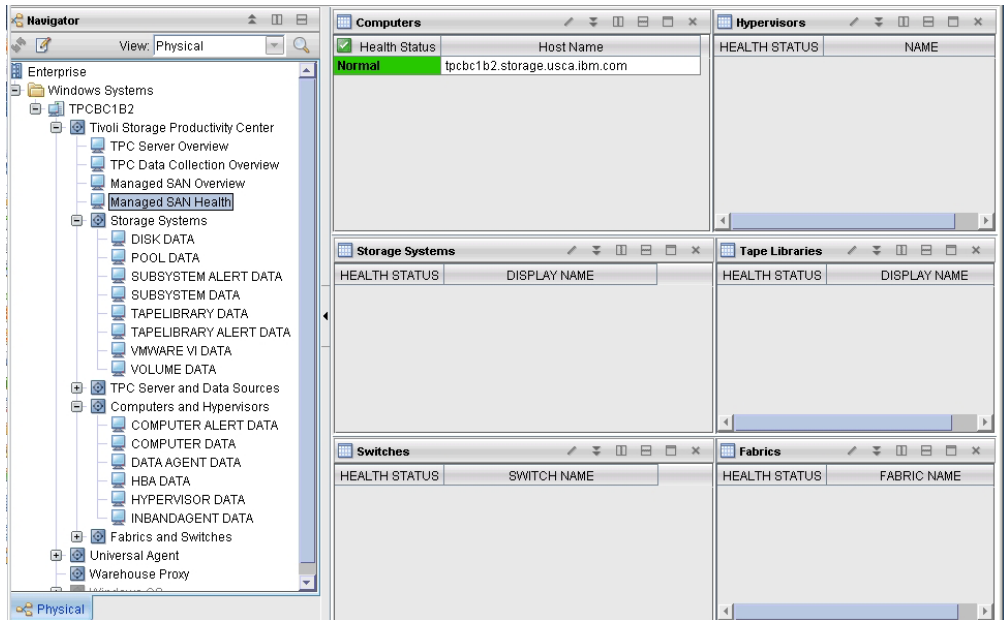


Figure 9. Managed SAN Health report

This workspace shows the following information:

**Computers**

Displays a list of all managed computers and their health status.

**Hypervisors**

Displays a list of all managed VMware hypervisors and their health status.

**Storage Systems**

Displays a list of all managed storage systems and their health status.

**Fabrics**

Displays a list of all managed fabrics and their health status.

**Switches**

Displays a list of all managed switches and their health status.

**Detailed Tivoli Storage Productivity Center reports**

A number of reports provide detailed information about the resources that are monitored by Tivoli Storage Productivity Center. To access these reports, click any of the following resources in the navigation tree of the Navigator window:

**Storage Systems**

Under **Storage Systems**, you can view these detailed reports:

- DISK DATA
- POOL DATA
- SUBSYSTEM ALERT DATA
- SUBSYSTEM DATA
- VMWARE VI DATA
- VOLUME DATA

**TPC Server and Data Sources**

Under TPC Server and Data Sources, you can view these detailed reports:

- CIMOM DATA

- DATASERVER API
- DEVICESERVER API
- JOB STATUS DATA
- REGISTERED NATIVE API DATA
- OTHER DEVICE DATA
- PERFORMANCE MONITOR DATA
- PROBE STATUS DATA
- SCHEDULED JOB DATA
- TPC SERVER DATA

#### **Computers and Hypervisors**

Under **Computers and Hypervisors**, you can view these detailed reports:

- COMPUTER ALERT DATA
- COMPUTER DATA
- HBA DATA
- HYPERVISOR DATA

#### **Fabrics and Switches**

Under **Fabrics and Switches**, you can view these detailed reports:

- FABRIC DATA
- FABRIC ALERT DATA
- OUTBANDAGENT DATA
- SWITCH DATA
- SWITCH ALERT DATA





---

## Appendix A. Reference

View reference information that is related to Tivoli Storage Productivity Center. Topics include information about alerts, fabrics, commands, configuration and log files, performance metrics, protocols, standards, and accessibility features.

---

### Alerts

Determining when and how you are alerted to conditions and violations of performance thresholds within a storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Alerts are triggered by conditions and violations of performance thresholds that are detected during data collection and event processing. For storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system, events are polled every minute from the resource. For the GPFS/GSS storage system, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.

#### Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some triggering conditions, like performance thresholds, require you to enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- A performance threshold was violated
- A data collection job did not complete
- A change occurred in the storage infrastructure

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over utilized or under utilized. Tivoli Storage Productivity Center provides many metrics for measuring performance and determining violations of the thresholds that you specify.

#### Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, Tivoli Enterprise Console events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

**Acknowledging alerts:** Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate

that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

## Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable Tivoli Storage Productivity Center to listen to SNMP traps. The default port number is 162, and the default community is public.

### Related concepts:

“Alerting” on page 159

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines resources for the conditions and performance thresholds that you specify and automatically notifies you when changes or violations are detected.

### Related tasks:

“Acknowledging alerts” on page 184

Some alerts in Tivoli Storage Productivity Center are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

## Triggering conditions for alerts

Define alerts so that IBM Tivoli Storage Productivity Center automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

### Related concepts:

“Alerting” on page 159

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines resources for the conditions and performance thresholds that you specify and automatically notifies you when changes or violations are detected.

### Related tasks:

“Adding resources” on page 127

You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

## Triggering conditions for storage system alerts

Triggering conditions are events in a storage system that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a storage system.

Alerts for storage systems can be triggered for the following types of conditions:

- Violations of performance thresholds
- Configuration changes on the following resources:
  - Storage systems
  - Back-end storage systems
  - Disks
  - Filesets
  - File systems
  - GPFS for NAS devices
  - Network Shared Disks (NSDs)
  - Nodes
  - Pools
  - Shares
  - Volumes

### Tips:

- The type of storage system determines which triggering conditions are available. For example, triggering conditions for shares are available only for storage systems that are configured for file storage, such as Storwize V7000 Unified.
- Use the error message identifiers that are associated with triggering conditions to locate any additional information about the error messages that are related to triggering conditions. For more information about ALR error messages, see [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.messages.doc/messages.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.messages.doc/messages.html).

Table 111. Triggering conditions for configuration changes on storage systems

Triggering Condition	Explanation	Related Error Message
Cache Client Discovered <sup>1</sup>	A cache client is detected by a probe of the storage system.	ALR4403I
Cache Client Missing <sup>1</sup>	A previously monitored cache client is not detected by a probe.	ALR4404I
Home System Discovered <sup>1</sup>	A home system is detected on a fileset by a probe of the storage system.	ALR4393I
Home System Missing <sup>1</sup>	A previously monitored home system is not detected by a probe.	ALR4395W
Home System Removed <sup>1</sup>	A home system is removed from a fileset.	ALR4394I
Performance Monitor Failed	A performance monitor did not complete when it attempted to collect performance data about the storage system.	
Probe Failed	A probe of a storage system did not complete.	
Quota Discovered <sup>1</sup>	A quota is discovered by a probe of the storage system.	ALR4496I
Quota Exceeded <sup>1</sup>	A quota was exceeded.	ALR4482W
Storage System Missing <sup>2</sup>	A previously monitored storage system is not found during a probe.	ALR0048W

Table 111. Triggering conditions for configuration changes on storage systems (continued)

Triggering Condition	Explanation	Related Error Message
Storage System Status Change Offline	A storage system goes offline.	ALR4241E
Storage System Status Change Online	A storage system comes online.	ALR4242W
Storage System Version Change	A storage system version changes.	ALR4243W
<b>Notes:</b> <ul style="list-style-type: none"> <li><sup>1</sup> This alert applies only to storage systems that are configured for file storage, including Storwize V7000 Unified. This alert does not apply to GPFS/GSS.</li> <li><sup>2</sup> This condition does not trigger alerts for GPFS/GSS.</li> </ul>		

Table 112. Triggering conditions for configuration changes on back-end storage systems

Triggering Condition	Explanation	Related Error Message
Back-end Storage System Status Change Offline	A back-end storage system goes offline.	ALR4247W
Back-end Storage System Status Change Online	A back-end storage system comes online.	ALR4248I

Table 113. Triggering conditions for configuration changes on disks

Triggering Condition <sup>*</sup>	Explanation	Related Error Message
Disk Missing	A previously monitored disk is not found during a probe of the storage system.	ALR4325W
Disk Discovered	A disk is discovered during a probe of the storage system.	ALR4251W
<b>Note:</b> <sup>*</sup> These conditions do not trigger alerts for the following storage systems: <ul style="list-style-type: none"> <li>Storwize V3500</li> <li>Storwize V3700</li> <li>Storwize V5000</li> <li>Storwize V7000</li> <li>Storwize V7000 Unified</li> <li>SAN Volume Controller</li> </ul>		

Table 114. Triggering conditions for configuration changes on exports

Triggering Condition	Explanation	Related Error Message
Export Missing <sup>*</sup>	A previously monitored export is not found during a probe.	ALR4470W
Export Discovered	An export is discovered when a probe collects data about the storage system.	ALR4460I
<b>Note:</b> <sup>*</sup> This condition does not trigger alerts.		

Table 115. Triggering conditions for configuration changes on filesets.

Triggering Condition	Explanation	Related Error Message
Cache Fileset Discovered	A cache fileset is discovered by a probe of the storage system.	ALR4397I

Table 115. Triggering conditions for configuration changes on filesets (continued).

Triggering Condition	Explanation	Related Error Message
Cache Fileset Missing	A previously monitored cache fileset is not detected during a probe.	ALR4398W
Fileset Discovered*	A fileset is discovered by a probe of the storage system.	ALR4447I
Fileset Linked*	A fileset is linked to a file system.	ALR4448I
Fileset Missing*	A previously monitored fileset is not detected during a probe.	ALR4474W
Fileset Unlinked*	A fileset is unlinked from a file system.	ALR4455I
<b>Note:</b> * This alert applies only to storage systems that are configured for file storage, including Storwize V7000 Unified and GPFS/GSS.		

Table 116. Triggering conditions for configuration changes on file systems

Triggering Condition*	Explanation	Related Error Message
File System Low On Absolute Free Inodes <sup>1</sup>	A file system is low on free i-nodes based on an absolute value threshold.	ALR4458W
File System Low On Absolute Free Space <sup>1</sup>	A file system is low on free space based on an absolute value threshold.	ALR4433W
File System Low On Relative Free Inodes <sup>1</sup>	A file system is low on free i-nodes based on a percentage value threshold.	ALR4458W
File System Low On Relative Free Space <sup>1</sup>	A file system is low on free space based on a percentage value threshold.	ALR4433W
File System Mount State Changed to Error	The mount state for a file system reached error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified or IBM SONAS.	ALR4356E
Storage File System Discovered <sup>2</sup>	A file system was detected for the first time.	ALR0005I, ALR0005W
Storage File System Missing <sup>2</sup>	A file system is removed or unmounted from a cluster.	ALR0006E
<b>Notes:</b> <ul style="list-style-type: none"> <li>• <sup>1</sup> This condition triggers alerts for Storwize V7000 Unified, GPFS/GSS, and IBM SONAS.</li> <li>• <sup>2</sup>This alert applies to file systems on the following resources: <ul style="list-style-type: none"> <li>– Storage systems that are configured for file storage, including Storwize V7000 Unified and GPFS/GSS</li> <li>– Servers that are managed by Storage Resource agents</li> </ul> </li> </ul>		

Table 117. Triggering conditions for configuration changes on the GPFS for NAS devices

Triggering Condition	Explanation	Related Error Message
GPFS Status Reached Error Level	The General Parallel File System (GPFS) daemon is down for one of the nodes of the NAS device. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4370E
GPFS Status Reached Warning Level	One of the GPFS nodes cannot be reached, or a node is trying to form a quorum with the other available nodes. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4369W

Table 118. Triggering conditions for configuration changes on nodes

Triggering Condition	Explanation	Related Error Message
Memory Usage Reached Error Level	Memory use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4362E
Node Discovered - Block Storage, Node Discovered - File Storage	A node was discovered during a probe of the storage system.	ALR1246W
Node Status Change Offline - Block Storage	A node goes offline.	ALR4226W
Node Status Change Offline - File Storage	For GPFS/GSS, this condition triggers alerts when the status of a node changes to “down”, or one of the nodes cannot be reached.	ALR4226W
Node Status Change Online - Block Storage	A node comes online.	ALR4227I
Node Status Change Online - File Storage	For GPFS/GSS, this condition triggers alerts when the status of a node changes to “active”.	ALR4227I
Node Status Changed to Error	The status of a clustered CIFS for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4367E
Node Status Changed to Warning	The clustered CIFS for a NAS device is disabled. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4366W
Processor Usage Reached Error Level	Processor use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4359E
Processor Usage Reached Warning Level	Processor use for a NAS device reached the warning level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4360W

Table 119. Triggering conditions for configuration changes on NSDs

Triggering Condition	Explanation	Related Error Message
NSD Discovered	A Network Shared Disk (NSD) is discovered during a probe of the storage system.	ALR4503I
NSD Missing	A previously monitored NSD is not found during a probe.	ALR4505W
NSD State Change	The state of an NSD changed.	ALR4507W
NSD Status Reached Error Level	NSD use for a NAS device reached the error level.	ALR4385E
NSD Status Reached Normal Level	NSD use for a NAS device reached the normal level.	ALR4387I
NSD Status Reached Warning Level	NSD use for a NAS device reached the warning level.	ALR4386W

Table 120. Triggering conditions for configuration changes on pools

Triggering Condition	Explanation	Related Error Message
Pool Available Space Too Low*	The available space in a pool that is not reserved for volumes is too low. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted. You can specify two threshold boundary values, a warning and a critical boundary value, in gigabytes or as a percentage.	ALR4541E, ALR4542W
Pool Discovered	A pool is discovered during a probe of the storage system.	ALR4253I
Pool Shortfall Too High*	The percentage of remaining unallocated volume space in a pool that is not available to be allocated is too high. The higher the percentage, the more critical the shortfall of pool space. You can specify two threshold boundary values, a warning and a critical boundary value, as a percentage.	ALR4545E, ALR4546W
Pool Status Change	A pool is missing or rediscovered.	ALR4252W
Pool Status Change Offline	A pool goes offline.	ALR4254W
Pool Status Change Online	A pool comes online.	ALR4255I
Pool Virtual Allocation Too High*	The pool virtual allocation exceeds the warning or critical threshold boundary value. The pool virtual allocation is the amount of physical space in a pool that was committed to the total virtual capacity of the volumes in the pool. The virtual allocation can exceed 100% if the pool contains thin provisioned volumes and is over committed. You can specify two threshold boundary values, a warning and a critical boundary value, in gigabytes or as a percentage.	ALR4543E, ALR4544W
<p><b>Note:</b> * This condition triggers alerts only for pools that contain thin-provisioned volumes on the following storage systems:</p> <ul style="list-style-type: none"> <li>• DS8000</li> <li>• Storwize V3500</li> <li>• Storwize V3700</li> <li>• Storwize V5000</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> <li>• Hitachi VSP</li> <li>• FlashSystem V840</li> </ul>		

Table 121. Triggering conditions for configuration changes on volumes

Triggering Condition*	Explanation	Related Error Message
Volume Discovered	A new volume is discovered during a probe of a storage system.	ALR4325W
Volume Missing	A previously monitored volume is not found during a probe of the storage system.	ALR4326W
Volume Status Change Offline	A volume goes offline.	ALR4249W
Volume Status Change Online	A volume comes online.	ALR4250I

Table 122. Triggering conditions for performance thresholds

Triggering Condition	Explanation	Related Error Message
<p>Back-end Read Queue Time Threshold</p>	<p>The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources. Specify a threshold to receive alerts about potential performance problems on an MDisk. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk is too high.</p> <p><b>Tip:</b> A queue algorithm on a storage system determines the number of concurrent I/O operations that are sent to an MDisk. If there is any queuing (other than during a backup process), then performance might be improved if you resolve the queuing issue.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 3 milliseconds per operation</li> <li>• Critical boundary: greater than 5 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount.</p>	<p>ALR0568E, ALR0569W, ALR0570W, ALR0571E</p>
<p>Back-end Write Queue Time Threshold</p>	<p>The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources. Specify a threshold to receive alerts about potential performance problems on an MDisk. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk is too high.</p> <p><b>Tip:</b> A queue algorithm on a storage system determines the number of concurrent I/O operations that are sent to an MDisk. If there is any queuing (other than during a backup process), then performance might be improved if you resolve the queuing issue.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 3 milliseconds per operation</li> <li>• Critical boundary: greater than 5 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount.</p>	<p>ALR0572E, ALR0573W, ALR0574W, ALR0575E</p>



Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Back-end Read Response Time Threshold	<p>The average number of milliseconds for the back-end storage resources to respond to a read operation. For SAN Volume Controller models, this value represents the external response time of the managed disks.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system MDisk or array. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk or array is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 25 milliseconds per operation</li> <li>• Critical boundary: greater than 35 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount.</p>	ALR0560E, ALR0561W, ALR0562W, ALR0563E
Back-end Write Response Time Threshold	<p>The average number of milliseconds for the back-end storage resources to respond to a write operation. For SAN Volume Controller models, this value represents the external response time of the managed disks.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system MDisk or array. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk or array is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 80 milliseconds per operation</li> <li>• Critical boundary: greater than 120 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Write I/O Rate is less than a specified amount.</p>	ALR0564E, ALR0565W, ALR0566W, ALR0567E
Cache Holding Time Threshold	<p>The average number of seconds that I/O data for a storage system controller is held in the cache. A short cache-holding time indicates adverse performance.</p> <p>Specify a threshold to receive alerts about potential performance problems on a controller. Slow response times might occur if the amount of cache that is installed on a controller is insufficient to manage its workload, or if the workload on a controller is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: less than 60 seconds</li> <li>• Critical boundary: less than 30 seconds</li> </ul>	ALR0528E, ALR0529W, ALR0530W, ALR0531E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
CPU Utilization Threshold	<p>The average percentage of time that the processors on nodes are busy. Specify a threshold to receive alerts about potential performance problems on the nodes in a storage system. Slow response times might occur if a node is experiencing a hardware or software-related problem that causes an increase in CPU usage, or if the workload on the node is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0556E, ALR0557W, ALR0558W, ALR0559E
CRC Error Rate Threshold	<p>The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. CRC errors suggest accidental modification of a frame's data during transit, and can be used as an indicator of the relative quality of a port's connection.</p> <p>Specify a threshold to receive alerts about conditions on a port that might slow the performance of the resources to which those ports are connected. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than 0.33 counts per second</li> </ul>	ALR0616E, ALR0617W, ALR0618W, ALR0619E
Credit Recovery Link Reset Rate Threshold	<p>The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a port that might slow the performance of the resources to which those ports are connected. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .02 counts per second</li> <li>• Critical boundary: greater than 0.04 counts per second</li> </ul>	ALR0660E, ALR0661W, ALR0662W, ALR0663E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Disk Utilization Percentage Threshold	<p>The average percentage of time that the disks that are associated with an array are busy. Specify a threshold to receive alerts about potential performance problems on a storage system array. Slow response times might occur if an array is experiencing a hardware or software-related problem, or if the workload on the array is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 50%</li> <li>• Critical boundary: greater than 80%</li> </ul> <p><b>Tips:</b></p> <ul style="list-style-type: none"> <li>• For DS6000 and DS8000 storage systems, this threshold applies only to those ranks that are the only ranks in their associated extent pool.</li> <li>• Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern. If this situation occurs for your workload, consider enabling the Sequential I/O Percentage filter. Use this filter to ignore any violations of the Disk Utilization Percentage threshold for highly sequential workloads.</li> </ul>	ALR0500E, ALR0501W, ALR0502W, ALR0503E
Error Frame Rate Threshold	<p>The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol. Specify a threshold to receive alerts about conditions on a switch port or storage system port that might slow the performance of the resources to which those ports are connected.</p>	ALR0544E, ALR0545W, ALR0546W, ALR0547E
Extreme I/O Concurrency Percentage Threshold	<p>The average percentage of I/O operations for which the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for ports on DS8000 cannot exceed 2000.</p> <p>Specify a threshold to receive alerts when the physical capability of a port on the DS8000 is being exceeded. If a port reaches its maximum number of open exchanges, it responds to any additional I/O requests with a 'Busy' status. When a "Busy" status is returned by the port, typically the host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.</p>	ALR0696E, ALR0697W, ALR0698W, ALR0699E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
I/O Busy Percentage Threshold	<p>The average percentage of I/O operations or exchanges for which the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of I/O operations exceeds an internal DS8000 threshold.</p> <p>Specify a threshold to receive alerts when the physical capability of a port on the DS8000 is being exceeded. When a "Busy" status is returned by a port, typically the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.</p>	ALR0700E, ALR0701W. ALR0702W, ALR0703E
I/O Overrun Percentage Threshold	<p>The average percentage of I/O operations that the port on DS8000 had to discard because the number of concurrent I/O operations or exchanges for the port exceeded 2000.</p> <p>Specify a threshold to receive alerts when the physical capability of a port on the DS8000 is being exceeded. When an I/O operation is discarded, typically the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.</p>	ALR0704E, ALR0705W. ALR0706W, ALR0707E
Invalid Link Transmission Rate Threshold	<p>The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss. Transmission words are primitive elements that are used by the Fibre Channel protocol for transmission of data. A frame consists of 10 - 537 transmission words.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected port or cable. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0692E, ALR0693W. ALR0694W, ALR0695E
Invalid Transmission Word Rate Threshold	<p>The average number of bit errors per second that are detected. Transmission words are primitive elements that are used by the Fibre Channel protocol for transmission of data. A frame consists of 10 - 537 transmission words.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, storage system port, for the connected port or cable. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0620E, ALR0621W, ALR0622W, ALR0623E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Link Failure Rate Threshold	<p>The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or storage system port, the connected port, or for fibre link or GBICs.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0548E, ALR0549W, ALR0550W, ALR0551E
Link Reset Received Rate Threshold	<p>The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected fabric.</p>	ALR0680E, ALR0681W, ALR0682W, ALR0683E
Loss of Signal Rate Threshold	<p>The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.</p> <p>Specify a threshold to receive alerts about potential hardware problems on a switch port, a storage system port, or for the connected port.</p> <p><b>Tip:</b> If the threshold that you set is violated, it might mean that the connected port was reset, perhaps because the host that contains the HBA was rebooted. Therefore, isolated errors can be ignored in most cases, but a rate that remains consistently high over time might need to be investigated.</p>	ALR0640E, ALR0641W, ALR0642W, ALR0643E
Loss of Sync Rate Threshold	<p>The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected port.</p>	ALR0636E, ALR0637W, ALR0638W, ALR0639E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Non-preferred Node Usage Percentage Threshold	<p>The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.</p> <p>Under normal conditions, the non-preferred node usage percentage should be zero for all I/O groups. Specify a threshold to receive alerts about the following issues:</p> <ul style="list-style-type: none"> <li>• Potential configuration problems for some of the host servers that use storage from the storage system.</li> <li>• Data paths from the host servers to the storage system through the fabric are not operational.</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Total I/O Rate of an I/O group is less than a specified amount.</p>	ALR0592E, ALR0593W, ALR0594W, ALR0595E
Overall Back-end Response Time Threshold	<p>The average number of milliseconds for the back-end storage resources to respond to a read or a write operation. For SAN Volume Controller models, this value represents the external response time of the managed disks.</p> <p>Specify a threshold to receive alerts about potential performance problems on an MDisk. Slow response times might occur if an MDisk is experiencing a hardware or software-related problem, or if the workload on the MDisk is too high.</p> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Total Back-end I/O Rate is less than a specified amount.</p>	ALR0512E, ALR0513W, ALR0514W, ALR0515E
Overall Port Response Time Threshold	<p>The average number of milliseconds to complete a send or receive operation. Specify a threshold to receive alerts about potential performance problems on a storage system port. Slow response times might occur if a port is experiencing a hardware or software-related problem, or if the workload on the port is too high.</p>	ALR0540E, ALR0541W, ALR0542W, ALR0543E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Peak Back-end Write Response Time Threshold	<p>The longest time for a back-end storage resource to respond to a write operation by a node. Specify a threshold to receive alerts about potential performance problems on the MDisks that are accessed by node. Slow response times might occur if an MDisk is experiencing a hardware or software-related problem, or if the workload on the MDisk is too high</p> <p>A violation of this threshold might mean that the SAN Volume Controller cache is setting a “partition-limit” for an MDisk group. The de-staged data from the cache for the MDisk group causes the cache to fill up. That is, writes are being received faster than they can be de-staged to disk. If delays reach 30 seconds or more, then the SAN Volume Controller switches to “short-term mode” where writes are no longer cached for the MDisk group.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 10,000 milliseconds</li> <li>• Critical boundary: greater than 30,000 milliseconds</li> </ul> <p><b>Tip:</b> Use this threshold only if the workload for resources in a storage system is guaranteed to always remain at a certain minimum level. Otherwise, many threshold violations might occur, even though there is really no underlying problem with the applications.</p>	ALR0596E, ALR0597W, ALR0598W, ALR0599E
Port Congestion Index Threshold	<p>The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 80 counts per second</li> <li>• Critical boundary: greater than 100 counts per second</li> </ul>	ALR0684E, ALR0685W, ALR0686W, ALR687E
Port Receive Bandwidth Percentage Threshold <sup>1</sup>	<p>The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port. Because fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%.</p> <p>Specify a threshold to receive alerts about potential performance problems on a port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75%</li> <li>• Critical boundary: greater than 85%</li> </ul>	ALR0612E, ALR0613W, ALR0614W, ALR0615E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Port Receive Utilization Percentage Threshold	<p>The average percentage of time that the port is busy receiving data. Specify a threshold to receive alerts about potential performance problems on a storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric.</p>	ALR0604E, ALR0605W, ALR0606W, ALR0607E
Port Send Bandwidth Percentage Threshold <sup>1</sup>	<p>The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port. Because fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%.</p> <p>Specify a threshold to receive alerts about potential performance problems on a switch port or storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75%</li> <li>• Critical boundary: greater than 85%</li> </ul>	ALR0608E, ALR0609W, ALR0610W, ALR0611E
Port Send Utilization Percentage Threshold	<p>The average percentage of time that the port is busy sending data. Specify a threshold to receive alerts about potential performance problems on a storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric.</p>	ALR0600E, ALR0601W, ALR0602W, ALR0603E
Port State Change Rate Threshold	<p>The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected. You can also use this threshold to help identify flapping ports.</p>	ALR0668E, ALR0669W, ALR0670W, ALR0671E
Port to Local Node Receive Queue Time Threshold	<p>The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that a node is waiting too long to receive data from other nodes and suggests that there is congestion around the ports on the fabric.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 1 millisecond per operation</li> <li>• Critical boundary: greater than 2 millisecond per operation</li> </ul>	ALR0588E, ALR0589W, ALR0590W, ALR0591E



Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Port to Local Node Send Queue Time Threshold	<p>The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that a node is waiting too long to send data to other nodes and suggests that there is congestion around the ports on the fabric.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 1 millisecond per operation</li> <li>• Critical boundary: greater than 2 milliseconds per operation</li> </ul>	ALR0584E, ALR0585W, ALR0586W, ALR0587E
Port to Local Node Receive Response Time Threshold	<p>The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that data is taking too long to be sent between nodes on the fabric, and suggests that there is either congestion around the FC ports, or there is an internal microcode problem on the SAN Volume Controller.</p> <p>This threshold has a default critical boundary of greater than 1 milliseconds per operation</p>	ALR0580E, ALR0581W, ALR0582W, ALR0583E
Port to Local Node Send Response Time Threshold	<p>The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that data is taking too long to be sent between nodes on the fabric, and suggests that there is either congestion around the FC ports, or there is an internal microcode problem on the SAN Volume Controller.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 1 millisecond per operation</li> <li>• Critical boundary: greater than 3 milliseconds per operation</li> </ul>	ALR0576E, ALR0577W, ALR0578W, ALR0579E
Primitive Sequence Protocol Error Rate Threshold	<p>The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port. Specify a threshold to receive alerts about conditions on a switch port or storage system that might slow the performance of the resources to which those ports are connected.</p>	ALR0672E, ALR0673W, ALR0674W, ALR0675E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Total Back-end Data Rate Threshold	<p>The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.</p> <p>Specify a threshold to receive alerts about storage system ranks, RAID arrays, MDisks, or pools (MDisk groups) that might have too high a workload.</p>	ALR0508E, ALR0509W, ALR0510W, ALR0511E
Total Back-end I/O Rate Threshold	<p>The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.</p> <p>Specify a threshold to receive alerts about storage system ranks, RAID arrays, MDisks, pools, controllers, nodes, or I/O groups that might have too high a workload.</p>	ALR0504E, ALR0505W, ALR0506W, ALR0507E
Total Data Rate Threshold	<p>The average number of MiB per second that are transferred for read operations and write operations. Specify a threshold to receive alerts about storage system volumes, RAID arrays, pools, controllers, nodes, I/O groups, or host connections that might have too high a workload.</p>	ALR0520E, ALR0521W, ALR0522W, ALR0523E
Total I/O Rate Threshold	<p>The average number of nonsequential I/O operations per second. This value includes both read and write operations. Specify a threshold to receive alerts about storage system volumes, RAID arrays, pools, controllers, or host connections that might have too high a workload.</p>	ALR0516E, ALR0517W, ALR0518W, ALR0519E
Total Port Data Rate Threshold	<p>The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations. Specify a threshold to receive alerts about switch ports or storage system ports that might have too high a workload.</p>	ALR0536E, ALR0537W, ALR0538W, ALR0539E
Total Port I/O Rate Threshold	<p>The average number of send operations and receive operations per second. Specify a threshold to receive alerts about storage system ports that might have too high a workload.</p>	ALR0532E, ALR0533W, ALR0534W, ALR0535E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Write-cache Delay Percentage Threshold	<p>The percentage of I/O operations that are delayed because there is not enough space in the write cache. The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system controller or node. For example, the following issues might be causing performance problems:</p> <ul style="list-style-type: none"> <li>• The amount of cache that is installed on the controller or node is insufficient to manage its workload.</li> <li>• The controller or node that is experiencing a hardware or software-related problem that causes cache slots to be occupied longer than necessary. This problem occurs when there is a drop in write performance, or the controller or node is being overwhelmed with too high a write workload.</li> </ul> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 3%</li> <li>• Critical boundary: greater than 10%</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount. The pre-populated filter value is 10 I/Os per second.</p>	ALR0524E, ALR0525W, ALR0526W, ALR0527E
Zero Buffer Credit Percentage Threshold <sup>3</sup>	<p>The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 25%</li> <li>• Critical boundary: greater than 66%</li> </ul>	ALR0628E, ALR0629W, ALR0630W, ALR0631E
Zero Buffer Credit Rate Threshold <sup>2</sup>	<p>The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75,000,000 milliseconds</li> <li>• Critical boundary: greater than 198,000,000 milliseconds</li> </ul>	ALR0648E, ALR0649W, ALR0650W, ALR0651E

Table 122. Triggering conditions for performance thresholds (continued)

Triggering Condition	Explanation	Related Error Message
Zero Buffer Credit Timer Threshold	The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.	ALR0624E, ALR0625W, ALR0626W, ALR0627E
Zero Receive Buffer Credit Percentage Threshold	The amount of time, as a percentage, that the port on DS8000 had depleted its receive buffer credits. That is, the percentage of time that the port on DS8000 had no credit to provide to the sending port.	ALR0708E, ALR0709W, ALR0710W, ALR0711E
Zero Send Buffer Credit Percentage Threshold	The amount of time, as a percentage, that the port on DS8000 had depleted its send buffer credits. That is, the percentage of time that the receiving port had no credit to provide to the port on DS8000.	ALR0712E, ALR0713W, ALR0714W, ALR0715E

**Notes:**

1. This condition does not trigger alerts for the DS6000.
2. This condition triggers alerts for SAN Volume Controller and Storwize system, port, node, and I/O groups.
3. This condition triggers alerts only for SAN Volume Controller and Storwize ports.

**Related tasks:**

“Tracking performance issues on a storage system” on page 110

You can be automatically notified of performance issues on storage system by defining alerts. Alerts can be triggered when the configuration of a storage system changes, or when it is not performing as expected. In this scenario, you define a performance threshold to be notified when the total I/O rate of a storage system falls outside an expected range.

**Related reference:**

“Performance thresholds” on page 1100

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

**Triggering conditions for hypervisor alerts**

Triggering conditions are events on a hypervisor that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a hypervisor.

Alerts for hypervisors can be triggered for the following types of conditions:

- Configuration changes on hypervisors
- Configuration changes on virtual machines
- Configuration changes on clusters
- Configuration changes on file systems
- Configuration changes on disks

**Tip:** Use the error message identifiers that are associated with triggering conditions to locate any additional information about the error messages that are related to triggering conditions. For more information about ALR error messages,

see  [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.messages.doc/messages.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.messages.doc/messages.html).

Table 123. Triggering conditions for configuration changes on hypervisors

Triggering Condition	Explanation	Related Error Message
Hypervisor Missing	A hypervisor that was previously discovered is missing.	ALR4321W
Probe Failed	The firmware version of host bus adapter changes.	ALR4035W
Server Status Change Offline	A monitored server goes offline.	ALR4273E, ALR4044W
Server Status Change Online	A monitored server comes online.	ALR4274I, ALR4045I

Table 124. Triggering conditions for configuration changes on virtual machines

Triggering Condition	Explanation	Related Error Message
Virtual Memory Increased	A new virtual machine was added.	ALR4316I
Virtual Memory Decreased	A virtual machine was removed.	ALR4317W

Table 125. Triggering conditions for configuration changes on clusters

Triggering Condition	Explanation	Related Error Message
Hypervisor Cluster Discovered	A cluster was discovered.	ALR4528I
Hypervisor Cluster Missing	A previously monitored hypervisor cluster is not found during a probe.	ALR4529I

Table 126. Triggering conditions for configuration changes on file systems

Triggering Condition	Explanation	Related Error Message
File System Low on Free Space	A file system is low on available storage space. This value is usually entered as a percent. This alert applies to file systems on servers that are managed by Storage Resource agents.	ALR0009W
File System Reconfigured	The physical space definition of a managed file system is reconfigured. This alert applies to file systems on the following resources: <ul style="list-style-type: none"> <li>Storage systems that are configured for file storage, including Storwize V7000 Unified and GPFS/GSS</li> <li>Servers that are managed by Storage Resource agents</li> </ul>	ALR0008I
File System Discovered	A file system was detected for the first time. This alert applies to file systems on the following resources: <ul style="list-style-type: none"> <li>Storage systems that are configured for file storage, including Storwize V7000 Unified and GPFS/GSS</li> <li>Servers that are managed by Storage Resource agents</li> </ul>	ALR0005I, ALR0005W

Table 127. Triggering conditions for configuration changes on disks

Triggering Condition	Explanation	Related Error Message
Disk Failure Predicted	A disk on a monitored hypervisor predicts that a disk failure is imminent.	ALR0007E
Disk Missing	A disk is removed from a monitored hypervisor.	ALR0004E
Grown Disk Defects Threshold Exceeded	The number of grown-disk defects exceeds a threshold that you specify.	ALR0012W

Table 127. Triggering conditions for configuration changes on disks (continued)

Triggering Condition	Explanation	Related Error Message
Disk Defect Discovered	A new grown disk defect is detected on a monitored disk.	ALR0011W
Disk Discovered	A new disk is discovered on a monitored hypervisor.	ALR0003I

### Triggering conditions for switch alerts

Triggering conditions are events on a switch that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a switch.

Alerts for switches can be triggered for the following types of conditions:

- Configuration changes on switches
- Configuration changes on blades
- Violations of performance thresholds

#### Tip:

- A performance monitor must collect data about a switch before Tivoli Storage Productivity Center can determine if performance thresholds are violated and an alert is generated.
- Use the error message identifiers that are associated with triggering conditions to locate any additional information about the error messages that are related to triggering conditions. For more information about ALR error messages, see [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.messages.doc/messages.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.messages.doc/messages.html).

Table 128. Triggering conditions for configuration changes on switches

Triggering Condition	Description	Related Error Message
Switch Status Degraded	The status of the switch changes from OK to a reduced level. This condition might occur if a fan or redundant power supply failed.	ALR4024W
Switch Version Change	The version of a switch changes.	ALR4023W

Table 129. Triggering conditions for configuration changes on blades

Triggering Condition	Description	Related Error Message
Switch Blade Change Offline	A switch module goes offline.	ALR4029E
Switch Blade Change Online	A switch module comes online.	ALR4030I

Table 130. Triggering conditions for performance thresholds

Triggering Condition*	Description	Related Error Message
Class-3 Receive Timeout Frame Rate Threshold	<p>The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a switch port that might slow the performance of the resources to which those ports are connected. This threshold is available only for Brocade switches.</p>	ALR0656E, ALR0657W, ALR0658W, ALR0659E
Class-3 Send Timeout Frame Rate Threshold	<p>The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a switch port that might slow the performance of the resources to which those ports are connected. This threshold is available only for Brocade switches.</p>	ALR0652E, ALR0652W, ALR0653W, ALR0654E
CRC Error Rate Threshold	<p>The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. CRC errors suggest accidental modification of a frame's data during transit, and can be used as an indicator of the relative quality of a port's connection.</p> <p>Specify a threshold to receive alerts about conditions on a port that might slow the performance of the resources to which those ports are connected. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than 0.33 counts per second</li> </ul>	ALR0616E, ALR0617W, ALR0618W, ALR0619E
Credit Recovery Link Reset Rate Threshold	<p>The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a port that might slow the performance of the resources to which those ports are connected. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .02 counts per second</li> <li>• Critical boundary: greater than 0.04 counts per second</li> </ul>	ALR0660E, ALR0661W, ALR0662W, ALR0663E

Table 130. Triggering conditions for performance thresholds (continued)

Triggering Condition *	Description	Related Error Message
Discarded Class-3 Frame Rate Threshold	The average number of class 3 frames per second that are discarded. Frames can be discarded due to many reasons, such as invalid addressing of header contents or a timeout that is caused by lack of buffer credits. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or for the connected fabric.	ALR0644E, ALR0645W, ALR0646W, ALR0647E
Discarded Frame Rate Threshold	The average number of frames per second that are discarded because host buffers are unavailable for the port. Frames can be discarded due to many reasons, such as invalid addressing of header contents or a timeout that is caused by lack of buffer credits. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or for the connected fabric.	ALR0632E, ALR0633W, ALR063W, ALR0635E
Error Frame Rate Threshold	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol. Specify a threshold to receive alerts about conditions on a switch port or storage system port that might slow the performance of the resources to which those ports are connected.	ALR0544E, ALR0545W, ALR0546W, ALR0547E
Invalid Link Transmission Rate Threshold	<p>The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss. Transmission words are primitive elements that are used by the Fibre Channel protocol for transmission of data. A frame consists of 10 - 537 transmission words.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected port or cable. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0692E, ALR0693W, ALR0694W, ALR0695E
Invalid Transmission Word Rate Threshold	<p>The average number of bit errors per second that are detected. Transmission words are primitive elements that are used by the Fibre Channel protocol for transmission of data. A frame consists of 10 - 537 transmission words.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, storage system port, for the connected port or cable. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0620E, ALR0621W, ALR0622W, ALR0623E



Table 130. Triggering conditions for performance thresholds (continued)

Triggering Condition *	Description	Related Error Message
Link Failure Rate Threshold	<p>The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or storage system port, the connected port, or for fibre link or GBICs.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>	ALR0548E, ALR0549W, ALR0550W, ALR0551E
Link Quality Percentage Threshold	<p>The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port. Specify a threshold to receive alerts about the estimated relative quality of the connection of a switch port or the connected port or cable.</p> <p>The default critical stress boundary for this threshold is set to greater than 75%.</p>	ALR0688E, ALR0689W, ALR0690W, ALR0691E
Link Reset Received Rate Threshold	<p>The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected fabric.</p>	ALR0680E, ALR0681W, ALR0682W, ALR0683E
Link Reset Transmitted Rate Threshold	<p>The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or for the connected fabric.</p>	ALR0676E, ALR0677W, ALR0678W, ALR0679E
Loss of Signal Rate Threshold	<p>The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.</p> <p>Specify a threshold to receive alerts about potential hardware problems on a switch port, a storage system port, or for the connected port.</p> <p><b>Tip:</b> If the threshold that you set is violated, it might mean that the connected port was reset, perhaps because the host that contains the HBA was rebooted. Therefore, isolated errors can be ignored in most cases, but a rate that remains consistently high over time might need to be investigated.</p>	ALR0640E, ALR0641W, ALR0642W, ALR0643E

Table 130. Triggering conditions for performance thresholds (continued)

Triggering Condition *	Description	Related Error Message
Loss of Sync Rate Threshold	<p>The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected port.</p>	ALR0636E, ALR0637W, ALR0638W, ALR0639E
Port Congestion Index Threshold	<p>The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 80 counts per second</li> <li>• Critical boundary: greater than 100 counts per second</li> </ul>	ALR0684E, ALR0685W, ALR0686W, ALR0687E
Port Receive Bandwidth Percentage Threshold	<p>The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port. Because fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%.</p> <p>Specify a threshold to receive alerts about potential performance problems on a port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75%</li> <li>• Critical boundary: greater than 85%</li> </ul>	ALR0612E, ALR0613W, ALR0614W, ALR0615E

Table 130. Triggering conditions for performance thresholds (continued)

Triggering Condition*	Description	Related Error Message
Port Send Bandwidth Percentage Threshold	<p>The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port. Because fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%.</p> <p>Specify a threshold to receive alerts about potential performance problems on a switch port or storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75%</li> <li>• Critical boundary: greater than 85%</li> </ul>	ALR0608E, ALR0609W, ALR0610W, ALR0611E
Port State Change Rate Threshold	The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected. You can also use this threshold to help identify flapping ports.	ALR0668E, ALR0669W, ALR0670W, ALR0671E
Primitive Sequence Protocol Error Rate Threshold	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port. Specify a threshold to receive alerts about conditions on a switch port or storage system that might slow the performance of the resources to which those ports are connected.	ALR0672E, ALR0673W, ALR0674W, ALR0675E
RDY Priority Override Rate Threshold	The average number of times per second during which the sending of R_RDY or VC_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected. This threshold is available only for Brocade switches.	ALR0664E, ALR0665W, ALR0666W, ALR0667E
Total Port Data Rate Threshold	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations. Specify a threshold to receive alerts about switch ports or storage system ports that might have too high a workload.	ALR0536E, ALR0537W, ALR0538W, ALR0539E
Total Port Frame Rate Threshold	The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port. Specify a threshold to receive alerts about switch ports or storage system ports that might have too high a workload.	ALR0552E, ALR0553W, ALR0554W, ALR0555E
<p><b>Tip:</b> Because the violations of isolated temporary errors can be ignored in most cases, you can select the alert notification option <b>Suppress alerts unless thresholds are violated more than once within the following length of time</b> for a threshold. This option can help reduce the number of unnecessary alerts.</p>		

## Triggering conditions for fabric alerts

Triggering conditions are events in a fabric that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected in a fabric.

Alerts for fabrics can be triggered when the following conditions are detected:

Table 131. Triggering conditions for configuration changes on fabrics

Triggering Condition	Explanation	Related Error Message
Fabric to Switch Change	A switch to fabric association is discovered, rediscovered, or missing.	ALR4066I, ALR4067W, ALR4068I
Probe Failed	A probe of a fabric did not complete.	

**Tip:** Use the error message identifiers that are associated with triggering conditions to locate any additional information about the error messages that are related to triggering conditions. For more information about ALR error messages, see [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.messages.doc/messages.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.messages.doc/messages.html).

## Triggering conditions for server alerts

Triggering conditions are events on a server that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a server.

Alerts for servers can be triggered for the following types of conditions:

- Configuration changes on servers
- Configuration changes on clusters
- Configuration changes on file systems
- Configuration changes on disks

**Tip:** Use the error message identifiers that are associated with triggering conditions to locate any additional information about the error messages that are related to triggering conditions. For more information about ALR error messages, see [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/com.ibm.tpc\\_V524.messages.doc/messages.html](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/com.ibm.tpc_V524.messages.doc/messages.html).

Table 132. Triggering conditions for configuration changes on servers

Triggering Condition	Explanation	Related Error Message
HBA Driver Version Change	The driver version of host bus adapter changes.	ALR4034W
HBA Firmware Version Change	The firmware version of host bus adapter changes.	ALR4035W
Probe Failed	The firmware version of host bus adapter changes.	ALR4035W
Server Status Change Offline	A monitored server goes offline.	ALR4273E, ALR4044W
Server Status Change Online	A monitored server comes online.	ALR4274I, ALR4045I

Table 133. Triggering conditions for configuration changes on clusters

Triggering Condition	Explanation	Related Error Message
Cluster Resource Group Added	A new cluster resource group is added.	ALR0055I
Cluster Resource Group Moved	A cluster resource group is moved.	ALR0057I
Cluster Resource Group Removed	A cluster resource group is removed.	ALR0056I

Table 134. Triggering conditions for configuration changes on file systems

Triggering Condition	Explanation	Related Error Message
File System Low on Free Space	A file system is low on available storage space. This value is usually entered as a percent. This alert applies to file systems on servers that are managed by Storage Resource agents.	ALR0009W
File System Low on Relative Free Inodes	A file system is low on free i-nodes based on a percentage value threshold. This condition triggers alerts for Storwize V7000 Unified, GPFS/GSS, and IBM SONAS.	ALR4458W
File System Missing	A file system is removed or unmounted from a cluster. This alert applies to file systems on the following resources: <ul style="list-style-type: none"> <li>Storage systems that are configured for file storage, including Storwize V7000 Unified and GPFS/GSS</li> <li>Servers that are managed by Storage Resource agents</li> </ul>	ALR0006E
File System Discovered	A file system was detected for the first time. This alert applies to file systems on the following resources: <ul style="list-style-type: none"> <li>Storage systems that are configured for file storage, including Storwize V7000 Unified and GPFS/GSS</li> <li>Servers that are managed by Storage Resource agents</li> </ul>	ALR0005I, ALR0005W

Table 135. Triggering conditions for configuration changes on disks

Triggering Condition	Explanation	Related Error Message
Disk Failure Predicted	A disk on a monitored server predicts that a disk failure is imminent.	ALR0007E
Disk Missing	A disk is removed from a monitored server.	ALR0004E
Grown Disk Defects Threshold Exceeded	The number of grown-disk defects exceeds a threshold that you specify.	ALR0012W
Disk Defect Discovered	A new grown disk defect is detected on a monitored disk.	ALR0011W
Disk Discovered	A new disk is discovered on a monitored server.	ALR0003I

### Alert suppression actions

If you selected a threshold as a triggering condition for alerts, you can specify the conditions under which alerts are triggered and choose whether you want to suppress repeating alerts. Alerts can be suppressed to avoid generating too many alert log entries or too many actions when the triggering condition occurs often. You can view suppressed alerts in the constraint violation reports.

You can define the following options that enable you to specify conditions that trigger and suppress alerts. If a threshold is not selected as a triggering condition, these options are not available.

**Trigger alerts for both critical and warning conditions**

Generates an alert upon the violation of either critical or warning threshold boundaries. This is the default.

**Trigger alerts for critical conditions only**

Generates alerts only upon violation of one of the critical threshold boundaries. Violation of a warning boundary creates an entry in the constraint violation report, but does not result in an entry in the alert log or an action being triggered.

**Trigger no alerts**

Does not generate an alert upon violation of any threshold boundaries. Creates entries only in the constraint violation report.

**Do not suppress repeating alerts**

Does not suppress any repeating alerts. This is the default.

**Suppress alerts unless the Triggering Condition has been violated continuously for a specified length of time**

Generates alerts only if the triggering condition has occurred continuously within the length of time specified in the **Length of time** field. Alerts for the first and any subsequent occurrences of the triggering condition within the specified time in minutes will be suppressed. At the point that there have been consecutive occurrences with the specified number of minutes, an alert is generated. When the specified suppression period has expired, the cycle starts again. Note that the timing for this feature is based on the IBM Tivoli Storage Productivity Center server clock rather than the various system clocks.

This option is useful for cases when a single occurrence of the triggering condition may be insignificant, but repeated occurrences can signal a potential problem.

**Suppress alerts if a repeat violation has occurred within a specified length of time after the initial violation of the Triggering Condition**


Generates alerts only for the first occurrence of the triggering condition. Alerts for repeated occurrences of the triggering condition within the length of time specified in the **Length of time** field are suppressed. When the specified suppression period has expired, the cycle starts again. Note that the timing for this feature is based on the IBM Tivoli Storage Productivity Center server clock rather than the various system clocks.

This option is useful for avoiding e-mail messages or similar disruptive alerts if the same triggering condition occurs repeatedly in successive sample passes. This option generally useful for all threshold types.

## **Triggering actions and notifications for alerts**

Specify how you are notified when alert conditions are detected on resources and any actions to take as a result of those alerts. These settings can be defined globally for a resource, or for specific alert definitions that apply to a resource.

Table 136. Locations where you can specify alert actions and notification settings.

Task	Location in GUI
Specify the notification settings for all the alert definitions that apply to a resource	Resource page and resource details page  For example, to specify the global notification settings for a storage system, go to the <b>Storage &gt; Storage Systems</b> page, right-click a storage system, and select <b>Edit Alert Notification Settings</b> . These notification settings are applied globally to all the alert definitions that are specified for the selected storage system. You can override the global settings by specifying different settings for each alert definition.
Specify alert actions and notification settings for a specific alert definition	Resource page and resource details page  For example, to specify the notification settings for an alert definition that applies to a storage system, go to the <b>Storage &gt; Storage Systems</b> page, right-click a storage system, and select <b>Edit Alert Definitions</b> . Click the envelope icon  next to the alert definition that you want to edit. The notification settings that you define override any global settings for the resource.

You can specify that the following actions are taken when alert conditions are detected on monitored resources:

#### Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. Learn more.

**Email** Send emails to specific email addresses when an alert is detected on a monitored resource.

#### Netcool or OMNIbus EIF Probe Server

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive Tivoli Storage Productivity Center alerts.

#### SNMP trap

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

#### Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

#### How scripts are run

After an alert is triggered, scripts can be run from the server where Tivoli Storage Productivity Center is installed or from any server where a Storage Resource agent is deployed. For some alert conditions, scripts can be specified to run on the server where the change was detected.

You must include the script files that you want to run in the /scripts directory on the server where Tivoli Storage Productivity Center is installed. The default path of /scripts is determined by the operating system of the server:

- Windows: C:\Program Files\IBM\TPC\Data\scripts
- UNIX or Linux: /opt/IBM/TPC/Data/scripts or /usr/IBM/TPC/Data/scripts

**Tip:** When you run a script against a NAS filer, the script is run from the Storage Resource agent that is assigned to the file system where the condition was detected.

## How script names are resolved

When you run scripts in response to alerts, the names of those scripts are resolved by using the following criteria:

- The operating system of the server where the product is installed.
- The script name that you specify on the **Script name** field when you define or edit an alert definition.

The operating system of the server or Storage Resource agent where you choose to run a script determines how that script that is run:

### UNIX, Linux operating systems

A Storage Resource agent that is deployed on a UNIX or Linux operating system does not run scripts that have an extension. If the specified script name contains an extension, the agent ignores the extension and searches for a file of the same name (without an extension). For example, if you enter backup.vbs in the **Script name** field, a UNIX or Linux agent searches for a file with the name backup. If the script exists, the first line of the script is read and the appropriate interpreter is used to run the script. If a file cannot be found, no action is taken.

### Windows operating systems

A Storage Resource agent that is deployed on a Windows operating system runs scripts that have an extension. The extension of a script file determines which interpreter is used to run the script. If you specify a script name that does not contain an extension, the agent searches for the file name in the following order (based on its extension): .com, .exe, .bat, .cmd, .vbs, .vbe, .js, .jse. If a file cannot be found, no action is taken.

**Tip:** You can run scripts with the same name across multiple operating systems. To run a script on both a UNIX or Linux and Windows operating systems, you must have two versions of that file in the /scripts directory. For example, if you want to run a provisioning script against Storage Resource agents on UNIX or Linux and Windows servers, you must have two versions of that file (provision.bat and provision) in /scripts.

## Alert display colors

In the stand-alone GUI, status symbols for alerts are shown in the colors that are listed for the alert action.

When you have active alerts that are shown within the nodes under the **IBM Tivoli Storage Productivity Center > Alerting > Alert Log** node in the navigation



tree, the indicator next to those nodes are displayed as green or red depending on their corresponding alerts. The following table shows the colors that are generated for the subnodes by their different alerts:

*Table 137. Alert colors*

<b>Alert</b>	<b>Color</b>
Probe failed	Red
RAM increase	Green
RAM decrease	Red
Memory increase	Green
Memory decrease	Red
New disk	Green
Missing disk	Red
Disk defect	Red
Disk failure predicted	Red
New file system	Green
File system reconfigured	Red
Scan failed	Red
Missing file system	Red
File system free space	Red
File system free Inodes	Red
Computer unreachable	Red
New directory	Red
Missing directory	Red
Directory quota violation	Red
File system quota violation	Red
Computer quota violation	Red
Network quota violation	Red
Constraint violation	Red
Failed job	Red
Disk defect quota	Red
Ping failed	Red
Discovery failed	Red
Aggregator failed	Red
File system quota failed	Red
Computer quota failed	Red
Network quota failed	Red
Computer discovered	Red
Report failed	Red
DB probe failed	Red
DB table space quota violation	Red
DB instance quota violation	Red
DB net quota violation	Red

Table 137. Alert colors (continued)

Alert	Color
DB table space quota failed	Red
DB instance quota failed	Red
DB net quota failed	Red
Scan failed	Red
Archive log overflow	Red
New table space	Red
Table space free space	Red
Table space offline	Red
Table space fragmented	Red
Table space extent size	Red
Dropped table space	Red
Table fragmented	Red
Table maximum extents	Red
Table size	Red
Table chained rows	Red
Table wasted space	Red
Dropped table	Red
NAS discovery failed	Red
NAS computer discovered	Green
Script failed	Red
NetApp quota job failed	Red
Device free space less	Red
Device free space greater	Red
Dropped device	Red
New device	Red
DB log free space	Red
DB dump time	Red
Disk array discovered	Green
Disk array missing	Red
Filer missing	Red
Cache increase	Green
Cache decrease	Red
Automatic file system extension	Red
Stopped automatic file system extension	Red
Extend job failed	Red
Upgrade failed	Red
Added cluster resource group	Green
Removed cluster resource group	Green
Moved cluster resource group	Green

Table 137. Alert colors (continued)

Alert	Color
DB report failed	Red
File backup failed	Red
Hypervisor discovered	Green
Hypervisor missing	Red
Virtual machine discovered	Green
Virtual machine missing	Red

## Performance data

Use the IBM Tivoli Storage Productivity Center to view the performance data of devices in your environment.

### Performance metrics

Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.

You can view performance metrics in the web-based GUI and stand-alone GUI at the following locations:

#### Resource pages in the web-based GUI

To access a view of performance metrics for a resource, go to the following pages in the web-based GUI:

- Server details page > Related Resources
- Storage systems page, Storage system details page
- Switches page, Switch details page

#### Reports in the web-based GUI

To view detailed performance reports for multiple resources in the web-based GUI, hover the mouse pointer over **Cognos** in the menu bar and select an option.

#### Reports in the stand-alone GUI

To view detailed performance reports for multiple resources in the stand-alone GUI, run the reports that are available under the following nodes in the navigation tree:

- **Disk Manager > Reporting**
- **Fabric Manager > Reporting**
- **Tivoli Storage Productivity Center > System Reports**

### Common information

The following table contains information that is common to performance reports in the stand-alone GUI.

Column	Description
Time	Date and time that the data was collected

Column	Description
Interval	<p>Size of the sample interval in seconds. You can specify a minimum interval length of 5 minutes and a maximum interval length of 60 minutes for the following storage systems:</p> <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server</li> <li>• DS6000</li> <li>• DS8000</li> <li>• XIV system</li> </ul> <p>For SAN Volume Controller models earlier than V4.1, you can specify a minimum interval length of 15 minutes and a maximum interval length of 60 minutes. For SAN Volume Controller models V4.1 and later, you can specify a minimum interval length of 5 minutes, and a maximum interval length of 60 minutes.</p>

## Metrics for ESS and DS6000 and DS8000 storage systems

When you view metrics for the ESS and DS6000 and DS8000 storage systems, you must take into account the following differences between IBM Tivoli Storage Productivity Center reports and the native reports of those systems:

- Tivoli Storage Productivity Center reports display port performance metrics as send and receive metrics, such as Send Data Rate and Receive Data Rate.
- Reports that are included with storage systems (for example, reports based on data that is collected by the DS CLI) display port performance metrics as read and write metrics, such as Bytread and Bytewrite.

When a host runs a read operation, the DS port sends data to the host. Therefore, "read" metrics in DS6000 and DS8000 reports correspond to "send" metrics in Tivoli Storage Productivity Center reports. When a host runs a write operation, DS6000 and DS8000 ports receive data from the host. Therefore, "write" metrics in DS6000 and DS8000 reports correspond to "receive" metrics in Tivoli Storage Productivity Center reports. When you view port Peer-to-Peer Remote Copy (PPRC) performance metrics, you must take into account the following additional differences between Tivoli Storage Productivity Center reports and native reports for storage systems:

- Metrics for PPRC reads in storage system native reports are represented as PPRC receives in Tivoli Storage Productivity Center (reads = receives).
- Metrics for PPRC writes in storage system native reports are represented as PPRC sends in Tivoli Storage Productivity Center (writes = sends).

**Restriction:** Tivoli Storage Productivity Center does not calculate volume-based metrics if thin-provisioned volumes are allocated in an extent pool that consists of multiple ranks. In this case, the columns for volume-based metrics display the value N/A in the Storage Subsystem Performance > By Array report for the arrays that are associated with that extent pool. However, if no thin-provisioned volumes are allocated in a multi-rank extent pool, or if the space efficient volumes are allocated in an extent pool that consists of a single rank, then this limitation does not apply and all volume-based metrics are displayed in the By Array report.

## Metrics for XIV systems

To distinguish between metrics that were introduced in different versions of the XIV systems, the following conventions are used:

- <sup>2</sup> is displayed next to metrics that are available in the XIV system version 10.2.2 or later.
- <sup>3</sup> is displayed next to metrics that are available in the XIV system version 10.2.4 or later.

For example:

- The Read I/O Rate (overall) metric is available for the XIV systems version 10.2.2 and later. In the Resources: components column of the list of metrics, the entry for Read I/O Rate (overall) is displayed like this example: XIV<sup>2</sup>: volume, module, storage system
- The Small Transfers Response Time metric is available for the XIV systems version 10.2.4 and later. In the Resources: components column of the list of metrics, the entry for Small Transfers Response Time is displayed in the following format: XIV<sup>3</sup>: volume, module, storage system

**Restriction:** XIV systems do not track performance statistics for volumes that were never used. Because there are no performance statistics, Tivoli Storage Productivity Center does not display metrics for these volumes and their related components.

## Metrics for SAN Volume Controller

Metrics that apply to SAN Volume Controller also apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified. For example, the Read I/O Rate (overall) metric is available for SAN Volume Controller and Storwize storage system volumes, even though only SAN Volume Controller is shown in the Resources: components column of the following lists.

### Volume metrics

The following metrics are available for the volumes in storage systems.

Metric	Resources: components <sup>1</sup>	Description
<b>I/O Rates</b>		
Read I/O Rate (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for nonsequential read operations for a component over a specified time interval.
Read I/O rate (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for sequential read operations for a component over a specified time interval.

Metric	Resources: components <sup>1</sup>	Description
Read I/O Rate (overall)	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller<sup>1</sup>: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of I/O operations per second for both sequential and nonsequential read operations for a component over a specified time interval.
Write I/O Rate (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for nonsequential write operations for a component over a specified time interval.
Write I/O Rate (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for sequential write operations for a component over a specified time interval.
Write I/O Rate (overall)	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of I/O operations per second for both sequential and nonsequential write operations for a component over a specified time interval.
Total I/O Rate (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for nonsequential read and write operations for a component over a specified time interval.

Metric	Resources: components <sup>1</sup>	Description
Total I/O Rate (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for sequential read and write operations for a component over a specified time interval.
Total I/O Rate (overall)	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of I/O operations per second for both sequential and nonsequential read and write operations for a component over a specified time interval.
Global Mirror Write I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of write operations per second issued to the Global Mirror secondary site for a component over a specified time interval.
Global Mirror Overlapping Write Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average percentage of write operations that are issued by the Global Mirror primary site and are serialized overlapping writes for a component over a specified time interval. For SAN Volume Controller 4.3.1 and later, some overlapping writes are processed in parallel (are not serialized) and are excluded. For earlier versions of SAN Volume Controller, all overlapping writes were serialized.
Global Mirror Overlapping Write I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of serialized overlapping write operations per second encountered by the Global Mirror primary site for a component over a specified time interval. For SAN Volume Controller 4.3.1 and later, some overlapping writes are processed in parallel (are not serialized) and are excluded. For earlier versions of SAN Volume Controller, all overlapping writes are serialized.
HPF Read I/O Rate	DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of read operations per second that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
HPF Write I/O Rate	DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of write operations per second that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
Total HPF I/O Rate	DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of read and write operations per second that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.

<b>Metric</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
HPF I/O Percentage	DS8000: volume, RAID array, pool, controller, host connection, storage system	The percentage of all I/O operations that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
PPRC Transfer Rate	ESS/DS6000/DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of track transfer operations per second for PPRC usage for a component over a specified time interval. This metric shows the activity for the source of the PPRC relationship, but shows no activity for the target.
Small Transfers I/O Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are <= 8 KB.
Medium Transfers I/O Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 8 KB and <= 64 KB.
Large Transfers I/O Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 64 KB and <= 512 KB.
Very Large Transfers I/O Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 512 KB.
<b>Cache hit percentages</b>		
Read Cache Hits Percentage (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for nonsequential read operations for a component over a specified time interval.
Read Cache Hits Percentage (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for sequential read operations for a component over a specified time interval.
Read Cache Hits Percentage (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system  SAN Volume Controller: volume, node, I/O group, host connection, storage system  SMI-S BSP: volume, controller, storage system  XIV <sup>1</sup> : volume, pool, module, host connection, storage system	Percentage of cache hits for both sequential and nonsequential read operations for a component over a specified time interval.



Metric	Resources: components <sup>1</sup>	Description
Write Cache Hits Percentage (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for nonsequential write operations for a component over a specified time interval.
Write Cache Hits Percentage (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for sequential write operations for a component over a specified time interval.
Write Cache Hits Percentage (overall)	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Percentage of cache hits for both sequential and nonsequential write operations for a component over a specified time interval.
Total Cache Hits Percentage (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for nonsequential read and write operations for a component over a specified time interval.
Total Cache Hits Percentage (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for sequential read and write operations for a component over a specified time interval.
Total Cache Hits Percentage (overall)	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV: volume, pool, module, host connection, storage system</p>	Percentage of cache hits for both sequential and nonsequential read and write operations for a component over a specified time interval.

<b>Metric</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
Readahead Percentage of Cache Hits	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of all read cache hits, which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of all write cache hits, which occurred on already dirty data in the cache.
Read Data Cache Hit Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of read data that was read from the cache over a specified time interval.
Write Data Cache Hit Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of write data that was written to the cache over a specified time interval.
Total Data Cache Hit Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	Percentage of all data that was read from or written to the cache for a component over a specified time interval.
SSD Read Cache Hit Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	<p>The percentage of all read operations for a component that accessed SSD cache memory during a specified time interval. This metric is only available for XIV systems that use SSD drives as an additional layer of caching.</p> <p>The value for this metric is also included in the value for the Read Cache Hit Percentage (overall) metric.</p>
SSD Read Data Cache Hit Percentage	XIV <sup>3</sup> : module, pool, volume, host connection, storage system	<p>The percentage of all read data for a component that was read from SSD cache memory during a specified time interval. This metric is only available for XIV systems that use SSD drives as an additional layer of caching.</p> <p>The value for this metric is also included in the value for the Read Data Cache Hit Percentage (overall) metric.</p>
<b>Data rates</b>		

Metric	Resources: components <sup>1</sup>	Description
Read Data Rate	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: controller, volume, storage system</p> <p>XIV<sup>2</sup>: module, pool, volume, host connection, storage system</p>	Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval.

<b>Metric</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
Small Transfers Data Percentage	XIV <sup>3</sup> : module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are <= 8 KB.
Medium Transfers Data Percentage	XIV <sup>3</sup> : module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 8 KB and <= 64 KB.
Large Transfers Data Percentage	XIV <sup>3</sup> : module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 64 KB and <= 512 KB.
Very Large Transfers Data	XIV <sup>3</sup> : module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 512 KB.
<b>Response times</b>		
Read Response Time	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.

Metric	Resources: components <sup>1</sup>	Description
Overall Response Time	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.
Peak Read Response Time	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The peak (worst) response time among all read operations.
Peak Write Response Time	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The peak (worst) response time among all write operations.
Global Mirror Write Secondary Lag	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The average number of additional milliseconds it takes to service each secondary write operation for Global Mirror, beyond the time that is required to service primary writes.
Overall Host Attributed Response Time Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The percentage of the average response time, both read response time and write response time, that can be attributed to delays from host systems. This metric is provided to help diagnose slow hosts and poorly performing fabrics. The value is based on the time that is taken for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes (for read) and the time that is taken for hosts to send the write data after the node responds to a transfer-ready notification (for write).
Read Cache Hit Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache hit operation over a specified time interval.
Write Cache Hit Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each write cache hit operation over a specified time interval.
Overall Cache Hit Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache hit operation and each write cache hit operation over a specified time interval.

<b>Metric</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
Read Cache Miss Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache miss operation over a specified time interval.
Write Cache Miss Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each write cache miss operation over a specified time interval.
Overall Cache Miss Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache miss operation and each write cache miss operation over a specified time interval.
Small Transfers Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are <= 8 KB.
Medium Transfers Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 8 KB and <= 64 KB.
Large Transfers Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 64 KB and <= 512 KB.
Very Large Transfers Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 512 KB.
SSD Read Cache Hit Response Time	XIV <sup>3</sup> : volume, pool, module, host connection, storage system	<p>Average number of milliseconds that was required to service each hit operation of SSD read memory for a component over a specified time interval. This metric is only available for XIV systems that use SSD drives as an additional layer of caching.</p> <p>The value for this metric is also included in the value for the Read Cache Hit Response Time metric.</p>
<b>Transfer sizes</b>		

Metric	Resources: components <sup>1</sup>	Description
Read Transfer Size	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of KB per I/O for read operations.
Write Transfer Size	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of KB per I/O for write operations.
Overall Transfer Size	<p>ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system</p> <p>SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system</p> <p>SMI-S BSP: volume, controller, storage system</p> <p>XIV<sup>2</sup>: volume, pool, module, host connection, storage system</p>	Average number of KB per I/O for read and write operations.
<b>Write-cache constraints</b>		

<b>Metric</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
Write-Cache Delay Percentage	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system  SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions for a component over a specified time interval. (The ratio of delayed operations to total I/Os.)
Write-Cache Delay I/O Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system  SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions for a component over a specified time interval.
Write-Cache Overflow Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The percentage of write operations that are delayed because there is not enough space in the write cache for a component over a specified time interval.
Write-Cache Overflow I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of write operations per second that are delayed because there is not enough space in the write cache for a component over a specified time interval.
Write-Cache Flush-through Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of tracks that is written to disk in flush-through mode for a component over a specified time interval.
Write-Cache Flush-through I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of tracks per second that are written to disk in flush-through mode for a component over a specified time interval.
Write-Cache Write-through Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of tracks that is written to disk in write-through mode for a component over a specified time interval.
Write-Cache Write-through I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of tracks per second that are written to disk in write-through mode for a component over a specified time interval.
<b>Record mode reads</b>		
Record Mode Read I/O Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection	Average number of I/O operations per second for record mode read operations for a component over a specified time interval.
Record Mode Read Cache Hit Percentage	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection	Percentage of cache hits for record mode read operations for a component over a specified time interval.
<b>Cache transfers</b>		



<b>Metric</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
Disk to Cache I/O Rate	ESS/DS6000/DS8000: volume, RAID array, pool, controller, host connection  SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of I/O operations (track transfers) per second for disk to cache transfers for a component over a specified time interval.
Cache to Disk I/O Rate	ESS/DS6000/DS8000: volume, RAID array, pool, controller, host connection  SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of I/O operations (track transfers) per second for cache to disk transfers for a component over a specified time interval.
<b>Miscellaneous computed values</b>		
Cache Holding Time	ESS, DS6000, DS8000: controller, storage system	Average cache holding time, in seconds, for I/O data in this storage system controller (cluster). Shorter time periods indicate adverse performance.
CPU Utilization	SAN Volume Controller: node, I/O group, storage system	Average utilization percentage of the processors.
Non-Preferred Node Usage Percentage	SAN Volume Controller: volume, I/O group, host connection	The overall percentage of I/O performed or data transferred through the non-preferred nodes of the volumes, for a component over a specified time interval.
Volume Utilization	ESS, DS6000, DS8000: volume  SAN Volume Controller: volume  SMI-S BSP: volume  XIV <sup>2</sup> : volume	The approximate utilization percentage of a volume over a specified time interval (the average percent of time that the volume was busy).
Pool Utilization	DS6000, DS8000: pool  SAN Volume Controller: pool  XIV <sup>2</sup> : pool	The approximate utilization percentage for the most active pools in the storage system. This utilization metric represents the average percent of time that a pool was busy. The value is based on an estimate of the workload for resources such as controllers, device adapters, volumes, and hard disks.
<p><b>Important:</b></p> <ul style="list-style-type: none"> <li><sup>1</sup> The "Resource: components" column shows the resources and components that are measured by a metric. Metrics that apply to SAN Volume Controller also apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.</li> <li><sup>2</sup> This metric is available in the XIV system version 10.2.2 or later.</li> <li><sup>3</sup> This metric is available in the XIV system version 10.2.4 or later.</li> </ul>		

## Disk metrics

The following table describes the metrics for the disks on storage systems.

Column	Resources: components <sup>1</sup>	Description
<b>I/O rates</b>		
Backend Read I/O Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller <sup>1</sup> : volume, pool, node, I/O group, storage system	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: volume, pool, node, I/O group, storage system	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: volume, pool, node, I/O group, storage system	Average number of I/O operations per second for read and write operations.
<b>Data rates</b>		
Backend Read Data Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for read operations.
Backend Write Data Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for write operations.

Column	Resources: components <sup>1</sup>	Description
Total Backend Data Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for read and write operations.
<b>Response times</b>		
Backend Read Response Time	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Backend Write Response Time	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Overall Backend Response Time	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Backend Read Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that each read operation was in the queue before it was issued to the back-end resource.
Backend Write Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that each write operation was in the queue before it was issued to the back-end resource.
Overall Backend Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that read and write operations were in the queue before they were issued to the back-end resource.
Peak Back-end Read Response Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The peak (worst) response time among all read operations for a component over a specified time interval. For SAN Volume Controller, it represents the external response time of the managed disks.

<b>Column</b>	<b>Resources: components<sup>1</sup></b>	<b>Description</b>
Peak Back-end Write Response Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The peak (worst) response time among all write operations for a component over a specified time interval. For SAN Volume Controller, it represents the external response time of the managed disks.
Peak Back-end Read Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The lower bound on the peak (worst) queue time for read operations for a component over a specified time interval. The queue time is the amount of time that the read operation was in the queue before it was issued to the back-end resource.
Peak Back-end Write Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The lower bound on the peak (worst) queue time for write operations for a component over a specified time interval. The queue time is the amount of time that the write operation was in the queue before it was issued to the back-end resource.
<b>Transfer sizes</b>		
Back-end Read Transfer Size	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of KB per I/O for read operations for a component over a specified time interval.
Back-end Write Transfer Size	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of KB per I/O for write operations for a component over a specified time interval.
Overall Back-end Transfer Size	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system  SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of KB per I/O for read and write operations for a component over a specified time interval.
<b>Disk utilization</b>		

Column	Resources: components <sup>1</sup>	Description
Disk Utilization Percentage	ESS, DS6000, DS8000: RAID array	The approximate utilization percentage of a rank over a specified time interval (the average percent of time that the disks associated with the array were busy). <b>Note:</b> Tivoli Storage Productivity Center does not calculate a value for this column if there are multiple ranks in the extent pool where thin-provisioned volumes are allocated. This column displays value of <b>N/A</b> for the reports in which it is displayed. However, if there is only a single rank in the extent pool, Tivoli Storage Productivity Center does calculate the value for this column regardless of the thin-provisioned volumes.
Sequential I/O Percentage	ESS, DS6000, DS8000: RAID array	Percentage of all I/O operations that were completed for a RAID array over a specified time interval that were sequential operations.
<b>Important:</b> <sup>1</sup> The "Resource: components" column shows the resources and components that are measured by a metric. Metrics that apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.		

## Port metrics

The following table describes the metrics for the ports on storage systems and switches.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Port Send I/O Rate	ESS, DS6000, DS8000: port, storage system  SAN Volume Controller <sup>1</sup> : port, node, I/O group, storage system  SMI-S BSP: port  XIV <sup>3</sup> : port	Average number of I/O operations per second for send operations for a port over a specified time interval.
Port Receive I/O Rate	ESS, DS6000, DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system  SMI-S BSP: port  XIV <sup>3</sup> : port	Average number of I/O operations per second for receive operations for a port over a specified time interval.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Total Port I/O Rate	ESS, DS6000, DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system  SMI-S BSP: port  XIV <sup>3</sup> : port	Average number of I/O operations per second for send and receive operations for a port over a specified time interval.
Port Send Frame Rate	switch port, switch	Average number of frames per second that are sent through the port.
Port Receive Frame Rate	switch port, switch	Average number of frames per second that are received by the port.
Total Port Frame Rate	switch port, switch	Average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.
Port to Host Send I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to host computers by a component over a specified time interval.
Port to Host Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from host computers by a component over a specified time interval.
Total Port to Host I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between host computers and a component over a specified time interval.
Port to Disk Send I/O Rate A	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to storage systems by a component over a specified time interval.
Port to Disk Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from storage systems by a component over a specified time interval.
Total Port to Disk I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between storage systems and a component over a specified time interval.
Port to Local Node Send I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Port to Local Node Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Total Port to Local Node I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between other nodes in the local SAN Volume Controller cluster and a component over a specified time interval.
Port to Remote Node Send I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.
Port to Remote Node Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from nodes in the remote SAN Volume Controller cluster.
Total Port to Remote Node I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between nodes in the remote SAN Volume Controller cluster and a component over a specified time interval.
Port FCP Send I/O Rate	ESS, DS6000, DS8000: port	Average number of send operations per second using the FCP protocol, for a port over a specified time interval.
Port FCP Receive I/O Rate	ESS, DS6000, DS8000: port	Average number of receive operations per second using the FCP protocol for a port over a specified time interval.
Total Port FCP I/O Rate	ESS, DS6000, DS8000: port	Average number of send and receive operations per second using the FCP protocol for a port over a specified time interval.
Port FICON Send I/O Rate	ESS, DS6000, DS8000: port	Average number of send operations per second using the FICON protocol for a port over a specified time interval.
Port FICON Receive I/O Rate	ESS, DS6000, DS8000: port	Average number of receive operations per second using the FICON protocol for a port over a specified time interval.
Total Port FICON I/O Rate	ESS, DS6000, DS8000: port	Average number of send and receive operations per second using the FICON protocol for a port over a specified time interval.
Port PPRC Send I/O Rate	ESS, DS6000, DS8000:: port, storage system	Average number of send operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Port PPRC Receive I/O Rate	ESS, DS6000, DS8000: port, storage system	Average number of receive operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Total Port PPRC I/O Rate	ESS, DS6000, DS8000: port, storage system	Average number of send and receive operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
<b>Data rates</b>		

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Port Send Data Rate	ESS, DS6000, DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system  SMI-S BSP: port  switch port, switch  XIV <sup>3</sup> : port	Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for send (read) operations for a port over a specified time interval.
Port Receive Data Rate	ESS, DS6000, DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system  SMI-S BSP: port  switch port, switch  XIV <sup>3</sup> : port	Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for receive (write) operations for a port over a specified time interval.
Total Port Data Rate	ESS, DS6000, DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system  SMI-S BSP: port  switch port, switch  XIV <sup>3</sup> : port	Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for send and receive operations for a port over a specified time interval.
Port Peak Send Data Rate	switch port	Peak number of megabytes (2 <sup>20</sup> bytes) per second that were sent by a port over a specified time interval
Port Peak Receive Data Rate	switch port	Peak number of megabytes (2 <sup>20</sup> bytes) per second that were received by a port over a specified time interval.
Port to Host Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second sent to host computers by a component over a specified time interval.
Port to Host Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second received from host computers by a component over a specified time interval.



Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Total Port to Host Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between host computers and a component over a specified time interval.
Port to Disk Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second sent to storage systems by a component over a specified time interval.
Port to Disk Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second received from storage systems by a component over a specified time interval.
Total Port to Disk Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between storage systems and a component over a specified time interval.
Port to Local Node Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second sent to other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Port to Local nNode Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second received from other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Total Port to Local Node Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between other nodes in the local SAN Volume Controller cluster and a component over a specified time interval.
Port to Remote Node Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second sent to nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.
Port to Remote Node Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second received from nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.
Total Port to Remote Node Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between nodes in the remote SAN Volume Controller cluster and a component over a specified time interval.
Port FCP Send Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2 <sup>20</sup> bytes) per second sent over the FCP protocol for a port over a specified time interval.
Port FCP Receive Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2 <sup>20</sup> bytes) per second received over the FCP protocol for a port over a specified time interval.
Total Port FCP Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2 <sup>20</sup> bytes) per second sent or received over the FCP protocol for a port over a specified time interval.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Port FICON Send Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2 <sup>^</sup> 20 bytes) per second sent over the FICON protocol for a port over a specified time interval.
Port FICON Receive Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2 <sup>^</sup> 20 bytes) per second received over the FICON protocol, for a port over a specified time interval.
Total Port FICON Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2 <sup>^</sup> 20 bytes) per second sent or received over the FICON protocol for a port over a specified time interval.
Port PPRC Send Data Rate	ESS, DS6000, DS8000: port, storage system	Average number of megabytes (2 <sup>^</sup> 20 bytes) per second sent for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Port PPRC Receive Data Rate	ESS, DS6000, DS8000: port, storage system	Average number of megabytes (2 <sup>^</sup> 20 bytes) per second received for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Total Port PPRC Data Rate	ESS, DS6000, DS8000: port, storage system	Average number of megabytes (2 <sup>^</sup> 20 bytes) per second transferred for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
<b>Response times</b>		
Port Send Response Time	ESS, DS6000, DS8000: port, storage system  XIV <sup>3</sup> : port	Average number of milliseconds that it took to service each send (read) operation for a port over a specified time interval.
Port Receive Response Time	ESS, DS6000, DS8000: port, storage system  XIV <sup>3</sup> : port	Average number of milliseconds that it took to service each receive (write) operation for a port over a specified time interval.
Overall Port Response Time	ESS, DS6000, DS8000: port, storage system  XIV <sup>3</sup> : port	Average number of milliseconds that it took to service each operation (send and receive) for a port over a specified time interval.
Port to Local Node Send Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send operation to another node in the local SAN Volume Controller cluster for a component over a specified time interval. This value is the external response time of the transfers.
Port to Local Node Receive Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each receive operation from another node in the local SAN Volume Controller cluster for a component over a specified time interval. This value is the external response time of the transfers.
Total Port to Local Node Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send or receive operation between another node in the local SAN Volume Controller cluster and a component over a specified time interval. This value is the external response time of the transfers.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Port to Local Node Send Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each send operation issued to another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port to Local Node Receive Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each receive operation from another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Total Port to Local Node Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each operation issued to another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port to Remote Node Send Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send operation to a node in the remote SAN Volume Controller cluster for a component over a specified time interval. This is the external response time of the transfers.
Port to Remote Node Receive Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each receive operation from a node in the remote SAN Volume Controller cluster for a component over a specified time interval. This value is the external response time of the transfers.
Total Port to Remote Node Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send or receive operation between a node in the remote SAN Volume Controller cluster and a component over a specified time interval. This value is the external response time of the transfers.
Port to Remote Node Send Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each send operation issued to a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port to Remote Node Receive Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each receive operation from a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Total Port to Remote Node Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each operation issued to a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port FCP Send Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all send operations over the FCP protocol for a port over a specified time interval.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Port FCP Receive Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all receive operations over the FCP protocol for a port over a specified time interval.
Overall Port FCP Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all I/O operations over the FCP protocol for a port over a specified time interval.
Port FICON Send Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all send operations over the FICON protocol for a port over a specified time interval.
Port FICON Receive Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all receive operations over the FICON protocol for a port over a specified time interval.
Overall Port FICON Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all I/O operations over the FICON protocol for a port over a specified time interval.
Port PPRC Send Response Time	ESS, DS6000, DS8000: port, storage system	Average number of milliseconds it took to service all send operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Port PPRC Receive Response Time	ESS, DS6000, DS8000: port, storage system	Average number of milliseconds it took to service all receive operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Overall Port PPRC Response Time	ESS, DS6000, DS8000: port, storage system	Average number of milliseconds it took to service all I/O operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
<b>Transfer sizes</b>		
Port Send Transfer Size	ESS, DS6000, DS8000: port, storage system SMI-S BSP: port	Average number of KB sent per I/O by a port over a specified time interval.
Port Receive Transfer Size	ESS, DS6000, DS8000: port, storage system SMI-S BSP: port	Average number of KB received per I/O by a port over a specified time interval.
Overall Port Transfer Size	ESS, DS6000, DS8000: port, storage system SMI-S BSP: port	Average number of KB transferred per I/O by a port over a specified time interval.
Port Send Frame Size	switch port, switch	The average size of a frame that is sent out through a port.
Port Receive Frame Size	switch port, switch	The average size of a frame that is received in through a port.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Overall Port Frame Size	switch port, switch	The average frame transfer size. This value includes frames that are sent and frames that are received by a port.
<b>Special computed values</b>		
Port Send Utilization Percentage	ESS, DS6000, DS8000: port	Average amount of time that the port was busy sending data over a specified time interval.
Port Receive Utilization Percentage	ESS, DS6000, DS8000: port	Average amount of time that the port was busy receiving data over a specified time interval.
Overall Port Utilization Percentage	ESS, DS6000, DS8000: port	Average amount of time that the port was busy sending or receiving data over a specified time interval.
Port Send Bandwidth Percentage	ESS, DS8000: port SAN Volume Controller: port switch, port XIV <sup>3</sup> : port	The approximate bandwidth utilization percentage for send operations by a port, based on its current negotiated speed.
Port Receive Bandwidth Percentage	ESS, DS8000: port SAN Volume Controller: port switch, port XIV <sup>3</sup> : port	The approximate bandwidth utilization percentage for receive operations by this port, based on its current negotiated speed.
Overall Port Bandwidth Percentage	ESS, DS8000: port SAN Volume Controller: port switch, port XIV <sup>3</sup> : port	The approximate bandwidth utilization percentage for send and receive operations by this port.
<b>Error rates</b>		
Error Frame Rate	switch port, switch DS8000: port, storage system	Average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
Discarded Frame Rate	switch port, switch	Average number of frames per second that are discarded because host buffers are unavailable for the port.
Link Failure Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	Number of link errors per second.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Loss of Sync Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	Average number of times per second that synchronization is lost. Synchronization is assumed lost after a timeout interval expires.
Loss of Signal Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	Average number of times per second that the signal is lost.
CRC Error Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	Average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.
Short Frame Rate	switch port, switch	Average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.
Long Frame Rate	switch port, switch	Average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.
Encoding Disparity Error Rate	switch port, switch	Average number of disparity errors per second that are received.
Discarded Class 3 Frame Rate	switch port, switch	Average number of class 3 frames per second that are discarded.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Class 3 Send Timeout Frame Rate	switch port, switch	Average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurred while waiting for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected. Available only for Brocade switches.
Class 3 Receive Timeout Frame Rate	switch port, switch	Average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurred while a transmitting port waited for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected. Available only for Brocade switches.
F-BSY Frame Rate	switch port, switch	Average number of F-BSY frames per second that are generated. An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.
F-RJT Frame Rate	switch port, switch	Average number of F-RJT frames per second that are generated. An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.
Primitive Sequence Protocol Error Rate	switch port, switch DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system	Average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Invalid Transmission Word Rate	switch port, switch DS8000: port, storage system  SAN Volume Controller: port, node, I/O group, storage system	Average number of bit errors per second that are detected. Errors include an 8b10 code violation in one or more characters, a K28.5 in the second, third, or fourth character positions, or an ordered set that had an incorrect Beginning Running Disparity.
Zero Buffer-Buffer Credit Timer	SAN Volume Controller: port, node, I/O group, storage system	Number of microseconds that the port was unable to send frames because of a lack of buffer credit since the last node reset.
Zero Buffer-Buffer Credit Percentage	SAN Volume Controller: port, node, I/O group, storage system	Percentage of sampling time that the port was unable to send frames because of a lack of buffer credit since the last node reset.

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Zero Buffer Credit Rate	switch port, switch	Average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected. Available only for Brocade switches.
Link Reset Transmitted Rate	switch port, switch DS8000: port, storage system	Average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Link Reset Received Rate	switch port, switch DS8000: port, storage system	Average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Credit Recovery Link Reset Rate	switch port, switch DS8000: port, storage system	Estimated average number of link resets per second that were performed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When troubleshooting a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected.
Out of Order Data Rate	DS8000: port, storage system	Average number of times per second that an out of order frame was detected for a port over a specified time interval.
Out of Order ACK Rate	DS8000: port, storage system	Average number of times per second that an out of order ACK frame was detected for a port over a specified time interval.
Duplicate Frame Rate	DS8000: port, storage system	Average number of times per second that a frame was received and also was detected as previously processed for a port over a specified time interval.
Invalid Relative Offset Rate	DS8000: port, storage system	Average number of times per second that a frame was received with invalid relative offset in the frame header for a port over a specific time interval.
Sequence Timeout Rate	DS8000: port, storage system	Average number of times per second the port detected a timeout condition on receiving sequence initiative for a Fibre Channel exchange for a port over a specified time interval.
RDY Priority Override Rate	switch port, switch	Average number of times per second that sending R_RDY or VC_RDY signals was a higher priority than sending frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected. Available only for Brocade switches.



Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
Port State Change Rate	switch port, switch	Average number of times per second that the state of a port changed to offline, online, or faulty. When you troubleshoot a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected. Available only for Brocade switches.
Port Congestion Index	switch port, switch SAN Volume Controller: port	Estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally between 0 (no congestion) to 100, but can exceed 100 if buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to view port conditions that might slow the performance of the resources to which those ports are connected.
Bad EOF CRC Error Rate	switch port	The average number of times per second that a cyclic redundancy check (CRC) error is detected in frames with a bad end-of-frame (EOF) indicator. In Brocade fabrics, a bad EOF in a frame can indicate that the frame has a known, previously detected CRC error. A good EOF in a frame with a CRC error indicates that the CRC error was not previously detected.
Link Quality Percentage	switch port	The estimated link quality of the switch port. The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the error statistics for the port.
Invalid Link Transmission Rate	switch port	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Extreme I/O Concurrency Rate	DS8000: port	The average number of times per second that the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for a port on DS8000 cannot exceed 2000.
Extreme I/O Concurrency Percentage	DS8000: port	The average percentage of I/O operations for which the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for ports on DS8000 cannot exceed 2000.
I/O Busy Rate	DS8000: port	The average number of times per second that the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of I/O operations or exchanges exceeds an internal DS8000 threshold.
I/O Busy Percentage	DS8000: port	The average percentage of I/O operations or exchanges for which the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of I/O operations exceeds an internal DS8000 threshold.

|  
|  
|  
|  
|  
|  
|  
|  
|  
|

Column	Resources: components <sup>1</sup>	Description
<b>I/O or frame rates</b>		
I/O Overrun Rate	DS8000: port	The average number of times per second that the port on DS8000 had to discard commands because the number of concurrent I/O operations or exchanges for the port exceeded 2000.
I/O Overrun Percentage	DS8000: port	The average percentage of I/O operations that the port on DS8000 had to discard because the number of concurrent I/O operations or exchanges for the port exceeded 2000.
Zero Send Buffer Credit Percentage	DS8000: port	The amount of time, as a percentage, that the port on DS8000 had depleted its send buffer credits. That is, the percentage of time that the receiving port had no credit to provide to the port on DS8000.
Zero Receive Buffer Credit Percentage	DS8000: port	The amount of time, as a percentage, that the port on DS8000 had depleted its receive buffer credits. That is, the percentage of time that the port on DS8000 had no credit to provide to the sending port.
<b>Important:</b> <ul style="list-style-type: none"> <li><sup>1</sup> The "Resource: components" column shows the resources and components that are measured by a metric. Metrics that apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.</li> <li><sup>2</sup> This metric is available in the XIV system version 10.2.2 or later.</li> <li><sup>3</sup> This metric is available in the XIV system version 10.2.4 or later.</li> </ul>		

#### Related concepts:

“Monitoring the performance of resources” on page 205

Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

## Performance thresholds

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

You can monitor the performance of your enterprise by creating alerts on performance thresholds for switches and storage subsystems. By creating alerts that are triggered by performance thresholds, you can be informed about performance issues in your enterprise. Threshold events tell you when a component has fallen outside of the user-defined values. For example, when a threshold value has reached critical stress.

### Threshold boundaries

You can establish your boundaries for the normal expected subsystem performance when defining storage subsystem alerts for performance threshold events. When the collected performance data samples fall outside the range you have set, you are notified of this threshold violation so you are aware of the potential problem. The

upper-threshold boundaries are **Critical Stress** and **Warning Stress**. The lower-threshold boundaries are **Warning Idle** and **Critical Idle**. Usually you will want the stress boundaries to be high numbers and the idle to be low numbers. The exception to this rule is **Cache Holding Time Threshold**, where you want the stress numbers to be low and the idle numbers to be high.

If you do not want to be notified of threshold violations for any boundaries, you can leave the boundary field blank and the performance data will not be checked against any value. For example, if you leave the **Critical Idle** and **Warning Idle** fields blank, no alerts will be sent for any idle conditions.

The **Ignore triggering condition when the sequential I/O percentage exceeds** check box is active only for the triggering condition **Disk Utilization Percentage Threshold**. It is a filter condition. The default is 80%.

The **Ignore triggering condition when the Back-end Read I/O Rate is less than** check box only applies to the **Back-end Read Response Time** and **Back-end Read Queue Time** thresholds.

The **Ignore triggering condition when the Back-end Write I/O Rate is less than** check box only applies to the **Back-end Write Response Time** and **Back-end Write Queue Time** thresholds.

The **Ignore triggering condition when the Total Back-end I/O Rate is less than** check box only applies to the **Overall Back-end Response Time** threshold.

The **Ignore triggering condition when the Total I/O Rate is less than** check box only applies to the **Non-preferred Node Usage Percentage** threshold.

The **Ignore triggering condition when the Write-cache Delay I/O Rate is less than** check box only applies to the **Write-cache Delay Percentage** threshold.

## Setting the thresholds

Only a few thresholds have defaults and on the other thresholds you will have to determine the best values for stress, idle, critical, and warning values so you can derive the maximum benefit without generating too many false alerts. Since suitable stress thresholds are highly dependant on the type of workload you are running, your exact hardware configuration, the number of physical disks, exact model numbers, and so forth, there are no easy or standard default rules.

One of the best approaches is to monitor your performance for a number of weeks and, using this historical data, determine reasonable values for each threshold setting. Once that is done you can fine tune these settings to minimize the number of false alerts.

### Related concepts:

“Alerting” on page 159

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines resources for the conditions and performance thresholds that you specify and automatically notifies you when changes or violations are detected.

“Monitoring the performance of resources” on page 205

Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

**Related tasks:**

“Adding resources” on page 127

You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

**Related reference:**

“Controller thresholds” on page 1110

Use thresholds for controllers to set limits on values such as data rates, I/O operations, and node usage, and to monitor cache-related performance.

“Port thresholds” on page 1112

Use thresholds for ports to set limits on values such as bandwidth utilization, data rates, and I/O operations.

“Array thresholds”

Use thresholds for arrays to set limits on values such as disk and CPU utilization, data rates, I/O operations, and response times.

**Array thresholds**

Use thresholds for arrays to set limits on values such as disk and CPU utilization, data rates, I/O operations, and response times.

**Tip:** Thresholds that are available for SAN Volume Controller are also available for Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.

Table 138. Array thresholds

Threshold (Metric)	Resource Type	Description
Back-end Read Queue Time	SAN Volume Controller MDisk	<p>The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources. Specify a threshold to receive alerts about potential performance problems on an MDisk. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk is too high.</p> <p><b>Tip:</b> A queue algorithm on a storage system determines the number of concurrent I/O operations that are sent to an MDisk. If there is any queuing (other than during a backup process), then performance might be improved if you resolve the queuing issue.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 3 milliseconds per operation</li> <li>• Critical boundary: greater than 5 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount.</p>

Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Back-end Write Queue Time	SAN Volume Controller MDisk	<p>The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources. Specify a threshold to receive alerts about potential performance problems on an MDisk. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk is too high.</p> <p><b>Tip:</b> A queue algorithm on a storage system determines the number of concurrent I/O operations that are sent to an MDisk. If there is any queuing (other than during a backup process), then performance might be improved if you resolve the queuing issue.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>Warning boundary: greater than 3 milliseconds per operation</li> <li>Critical boundary: greater than 5 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount.</p>
Back-end Read Response Time	SAN Volume Controller MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	<p>The average number of milliseconds for the back-end storage resources to respond to a read operation. For SAN Volume Controller models, this value represents the external response time of the managed disks.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system MDisk or array. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk or array is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>Warning boundary: greater than 25 milliseconds per operation</li> <li>Critical boundary: greater than 35 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount.</p>

Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Back-end Write Response Time	SAN Volume Controller MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	<p>The average number of milliseconds for the back-end storage resources to respond to a write operation. For SAN Volume Controller models, this value represents the external response time of the managed disks.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system MDisk or array. Slow response times might be caused by a hardware or software-related problem, or if the workload on an MDisk or array is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 80 milliseconds per operation</li> <li>• Critical boundary: greater than 120 milliseconds per operation</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Write I/O Rate is less than a specified amount.</p>
CPU Utilization Percentage	SAN Volume Controller node, I/O group, storage system	<p>The average percentage of time that the processors on nodes are busy. Specify a threshold to receive alerts about potential performance problems on the nodes in a storage system. Slow response times might occur if a node is experiencing a hardware or software-related problem that causes an increase in CPU usage, or if the workload on the node is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>

Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Disk Utilization Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000 array	<p>The average percentage of time that the disks that are associated with an array are busy. Specify a threshold to receive alerts about potential performance problems on a storage system array. Slow response times might occur if an array is experiencing a hardware or software-related problem, or if the workload on the array is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 50%</li> <li>• Critical boundary: greater than 80%</li> </ul> <p><b>Tips:</b></p> <ul style="list-style-type: none"> <li>• For DS6000 and DS8000 storage systems, this threshold applies only to those ranks that are the only ranks in their associated extent pool.</li> <li>• Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern. If this situation occurs for your workload, consider enabling the Sequential I/O Percentage filter. Use this filter to ignore any violations of the Disk Utilization Percentage threshold for highly sequential workloads.</li> </ul>
Overall Back-end Response Time	SAN Volume Controller MDisk	<p>The average number of milliseconds for the back-end storage resources to respond to a read or a write operation. For SAN Volume Controller models, this value represents the external response time of the managed disks.</p> <p>Specify a threshold to receive alerts about potential performance problems on an MDisk. Slow response times might occur if an MDisk is experiencing a hardware or software-related problem, or if the workload on the MDisk is too high.</p> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Total Back-end I/O Rate is less than a specified amount.</p>



Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Peak Back-end Write Response Time	SAN Volume Controller node	<p>The longest time for a back-end storage resource to respond to a write operation by a node. Specify a threshold to receive alerts about potential performance problems on the MDisks that are accessed by node. Slow response times might occur if an MDisk is experiencing a hardware or software-related problem, or if the workload on the MDisk is too high</p> <p>A violation of this threshold might mean that the SAN Volume Controller cache is setting a “partition-limit” for an MDisk group. The de-staged data from the cache for the MDisk group causes the cache to fill up. That is, writes are being received faster than they can be de-staged to disk. If delays reach 30 seconds or more, then the SAN Volume Controller switches to “short-term mode” where writes are no longer cached for the MDisk group.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 10,000 milliseconds</li> <li>• Critical boundary: greater than 30,000 milliseconds</li> </ul> <p><b>Tip:</b> Use this threshold only if the workload for resources in a storage system is guaranteed to always remain at a certain minimum level. Otherwise, many threshold violations might occur, even though there is really no underlying problem with the applications.</p>

Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Port to Local Node Receive Queue Time	SAN Volume Controller node	<p>The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that a node is waiting too long to receive data from other nodes and suggests that there is congestion around the ports on the fabric.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 1 millisecond per operation</li> <li>• Critical boundary: greater than 2 millisecond per operation</li> </ul>
Port to Local Node Receive Response Time	SAN Volume Controller node	<p>The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that data is taking too long to be sent between nodes on the fabric, and suggests that there is either congestion around the FC ports, or there is an internal microcode problem on the SAN Volume Controller.</p> <p>This threshold has a default critical boundary of greater than 1 milliseconds per operation</p>

Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Port to Local Node Send Queue Time	SAN Volume Controller node	<p>The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that a node is waiting too long to send data to other nodes and suggests that there is congestion around the ports on the fabric.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 1 millisecond per operation</li> <li>• Critical boundary: greater than 2 milliseconds per operation</li> </ul>
Port to Local Node Send Response Time	SAN Volume Controller node	<p>The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system node. A violation of this threshold means that data is taking too long to be sent between nodes on the fabric, and suggests that there is either congestion around the FC ports, or there is an internal microcode problem on the SAN Volume Controller.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 1 millisecond per operation</li> <li>• Critical boundary: greater than 3 milliseconds per operation</li> </ul>
Total Back-end Data Rate	SAN Volume Controller MDisk group and MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	<p>The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.</p> <p>Specify a threshold to receive alerts about storage system ranks, RAID arrays, MDisks, or pools (MDisk groups) that might have too high a workload.</p>

Table 138. Array thresholds (continued)

Threshold (Metric)	Resource Type	Description
Total Back-end I/O Rate	SAN Volume Controller MDisk group and MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	<p>The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.</p> <p>Specify a threshold to receive alerts about storage system ranks, RAID arrays, MDisks, pools, controllers, nodes, or I/O groups that might have too high a workload.</p>

## Controller thresholds

Use thresholds for controllers to set limits on values such as data rates, I/O operations, and node usage, and to monitor cache-related performance.

**Tip:** Thresholds that are available for SAN Volume Controller are also available for Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.

Table 139. Controller thresholds

Threshold (Metric)	Resource Type	Description
Cache Holding Time	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller	<p>The average number of seconds that I/O data for a storage system controller is held in the cache. A short cache-holding time indicates adverse performance.</p> <p>Specify a threshold to receive alerts about potential performance problems on a controller. Slow response times might occur if the amount of cache that is installed on a controller is insufficient to manage its workload, or if the workload on a controller is too high.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>Warning boundary: less than 60 seconds</li> <li>Critical boundary: less than 30 seconds</li> </ul>

Table 139. Controller thresholds (continued)

Threshold (Metric)	Resource Type	Description
Non-Preferred Node Usage Percentage	SAN Volume Controller I/O group	<p>The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.</p> <p>Under normal conditions, the non-preferred node usage percentage should be zero for all I/O groups. Specify a threshold to receive alerts about the following issues:</p> <ul style="list-style-type: none"> <li>• Potential configuration problems for some of the host servers that use storage from the storage system.</li> <li>• Data paths from the host servers to the storage system through the fabric are not operational.</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Total I/O Rate of an I/O group is less than a specified amount.</p>
Total Data Rate	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller, SAN Volume Controller I/O group	The average number of MiB per second that are transferred for read operations and write operations. Specify a threshold to receive alerts about storage system volumes, RAID arrays, pools, controllers, nodes, I/O groups, or host connections that might have too high a workload.
Total I/O Rate (overall)	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller	The average number of nonsequential I/O operations per second. This value includes both read and write operations. Specify a threshold to receive alerts about storage system volumes, RAID arrays, pools, controllers, or host connections that might have too high a workload.

Table 139. Controller thresholds (continued)

Threshold (Metric)	Resource Type	Description
Write-cache Delay Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller, SAN Volume Controller node	<p>The percentage of I/O operations that are delayed because there is not enough space in the write cache. The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.</p> <p>Specify a threshold to receive alerts about potential performance problems on a storage system controller or node. For example, the following issues might be causing performance problems:</p> <ul style="list-style-type: none"> <li>• The amount of cache that is installed on the controller or node is insufficient to manage its workload.</li> <li>• The controller or node that is experiencing a hardware or software-related problem that causes cache slots to be occupied longer than necessary. This problem occurs when there is a drop in write performance, or the controller or node is being overwhelmed with too high a write workload.</li> </ul> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 3%</li> <li>• Critical boundary: greater than 10%</li> </ul> <p>To reduce the number of unnecessary alerts, you can set a filter to ignore boundary violations if the Back-end Read I/O Rate is less than a specified amount. The pre-populated filter value is 10 I/Os per second.</p>

### Port thresholds

Use thresholds for ports to set limits on values such as bandwidth utilization, data rates, and I/O operations.

**Tip:** Thresholds that are available for SAN Volume Controller ports are also available for ports on Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.

Table 140. Port thresholds

Threshold (Metric)	Resource Type	Description
Class-3 Receive Timeout Frame Rate	Brocade switch port	<p>The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a switch port that might slow the performance of the resources to which those ports are connected. This threshold is available only for Brocade switches.</p>
Class-3 Send Timeout Frame Rate	Brocade switch port	<p>The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a switch port that might slow the performance of the resources to which those ports are connected. This threshold is available only for Brocade switches.</p>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
CRC Error Rate	Switch port, DS8000 port, SAN Volume Controller port	<p>The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. CRC errors suggest accidental modification of a frame's data during transit, and can be used as an indicator of the relative quality of a port's connection.</p> <p>Specify a threshold to receive alerts about conditions on a port that might slow the performance of the resources to which those ports are connected. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than 0.33 counts per second</li> </ul>



Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Credit Recovery Link Reset Rate	Switch port, DS8000 port	<p>The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>Specify a threshold to receive alerts about conditions on a port that might slow the performance of the resources to which those ports are connected. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .02 counts per second</li> <li>• Critical boundary: greater than 0.04 counts per second</li> </ul>
Discarded Class-3 Frame Rate	Switch port	<p>The average number of class 3 frames per second that are discarded. Frames can be discarded due to many reasons, such as invalid addressing of header contents or a timeout that is caused by lack of buffer credits. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or for the connected fabric.</p>
Discarded Frame Rate	Switch port	<p>The average number of frames per second that are discarded because host buffers are unavailable for the port. Frames can be discarded due to many reasons, such as invalid addressing of header contents or a timeout that is caused by lack of buffer credits. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or for the connected fabric.</p>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Error Frame Rate	Switch port, DS8000 port	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol. Specify a threshold to receive alerts about conditions on a switch port or storage system port that might slow the performance of the resources to which those ports are connected.
Extreme I/O Concurrency Percentage	DS8000 port	<p>The average percentage of I/O operations for which the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for ports on DS8000 cannot exceed 2000.</p> <p>Specify a threshold to receive alerts when the physical capability of a port on the DS8000 is being exceeded. If a port reaches its maximum number of open exchanges, it responds to any additional I/O requests with a 'Busy' status. When a "Busy" status is returned by the port, typically the host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.</p>
I/O Busy Percentage	DS8000 port	<p>The average percentage of I/O operations or exchanges for which the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of I/O operations exceeds an internal DS8000 threshold.</p> <p>Specify a threshold to receive alerts when the physical capability of a port on the DS8000 is being exceeded. When a "Busy" status is returned by a port, typically the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.</p>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
I/O Overrun Percentage	DS8000 port	<p>The average percentage of I/O operations that the port on DS8000 had to discard because the number of concurrent I/O operations or exchanges for the port exceeded 2000.</p> <p>Specify a threshold to receive alerts when the physical capability of a port on the DS8000 is being exceeded. When an I/O operation is discarded, typically the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.</p>
Invalid Link Transmission Rate	Switch port, DS8000 port	<p>The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss. Transmission words are primitive elements that are used by the Fibre Channel protocol for transmission of data. A frame consists of 10 - 537 transmission words.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected port or cable. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Invalid Transmission Word Rate	Switch port, DS8000 port, SAN Volume Controller port	<p>The average number of bit errors per second that are detected. Transmission words are primitive elements that are used by the Fibre Channel protocol for transmission of data. A frame consists of 10 - 537 transmission words.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, storage system port, for the connected port or cable. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>
Link Failure Rate	Switch port, DS8000 port, SAN Volume Controller port	<p>The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or storage system port, the connected port, or for fibre link or GBICs.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than .01 counts per second</li> <li>• Critical boundary: greater than .033 counts per second</li> </ul>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Link Quality Percentage Threshold	Switch port	<p>The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port. Specify a threshold to receive alerts about the estimated relative quality of the connection of a switch port or the connected port or cable.</p> <p>The default critical stress boundary for this threshold is set to greater than 75%.</p>
Link Reset Received Rate	Switch port, DS8000 port	<p>The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected fabric.</p>
Link Reset Transmitted Rate	Switch port, DS8000 port	<p>The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state. Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port or for the connected fabric.</p>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Loss of Signal Rate	Switch port, DS8000 port, SAN Volume Controller port	<p>The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.</p> <p>Specify a threshold to receive alerts about potential hardware problems on a switch port, a storage system port, or for the connected port.</p> <p><b>Tip:</b> If the threshold that you set is violated, it might mean that the connected port was reset, perhaps because the host that contains the HBA was rebooted. Therefore, isolated errors can be ignored in most cases, but a rate that remains consistently high over time might need to be investigated.</p>
Loss of Sync Rate	Switch port, DS8000 port, SAN Volume Controller port	<p>The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.</p> <p>Specify a threshold to receive alerts about potential hardware or configuration problems on a switch port, a storage system port, or for the connected port.</p>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Overall Port Response Time	TotalStorage Enterprise Storage Server port, DS6000 port, DS8000 port	<p>The average number of milliseconds to complete a send or receive operation. Specify a threshold to receive alerts about potential performance problems on a storage system port. Slow response times might occur if a port is experiencing a hardware or software-related problem, or if the workload on the port is too high.</p>
Port Congestion Index	Switch port, DS8000 port, SAN Volume Controller port, Storwize V7000 port	<p>The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 80 counts per second</li> <li>• Critical boundary: greater than 100 counts per second</li> </ul>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Port Receive Bandwidth Percentage	Switch port, TotalStorage Enterprise Storage Server port, DS8000 port, SAN Volume Controller port, XIV port	<p>The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port. Because fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%.</p> <p>Specify a threshold to receive alerts about potential performance problems on a port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75%</li> <li>• Critical boundary: greater than 85%</li> </ul>
Port Receive Utilization Percentage	TotalStorage Enterprise Storage Server port, DS6000 port, DS8000 port	<p>The average percentage of time that the port is busy receiving data. Specify a threshold to receive alerts about potential performance problems on a storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric.</p>



Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Port Send Bandwidth Percentage	Switch port, TotalStorage Enterprise Storage Server port, DS8000 port, SAN Volume Controller port, XIV port	<p>The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port. Because fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%.</p> <p>Specify a threshold to receive alerts about potential performance problems on a switch port or storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric. This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75%</li> <li>• Critical boundary: greater than 85%</li> </ul>
Port Send Utilization Percentage	TotalStorage Enterprise Storage Server port, DS6000 port, DS8000 port	<p>The average percentage of time that the port is busy sending data. Specify a threshold to receive alerts about potential performance problems on a storage system port. Slow response times might occur if the workload for a port is too high, if it is experiencing a hardware or software-related problem, or if there are congestion or configuration issues in the fabric.</p>
Port State Change Rate	Switch port, TotalStorage Enterprise Storage Server port, DS8000 port, SAN Volume Controller port, XIV port	<p>The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected. You can also use this threshold to help identify flapping ports.</p>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Primitive Sequence Protocol Error Rate	Switch port, DS8000 port, SAN Volume Controller port	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port. Specify a threshold to receive alerts about conditions on a switch port or storage system that might slow the performance of the resources to which those ports are connected.
RDY Priority Override Rate	Brocade switch port	The average number of times per second during which the sending of R_RDY or VC_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.  The threshold is available only for Brocade switches.
Total Port Data Rate	Switch port, TotalStorage Enterprise Storage Server, DS6000, DS8000port, XIV port	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations. Specify a threshold to receive alerts about switch ports or storage system ports that might have too high a workload.
Total Port Frame Rate	Switch port	The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port. Specify a threshold to receive alerts about switch ports or storage system ports that might have too high a workload.

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Total Port I/O Rate	Switch port, TotalStorage Enterprise Storage Server, DS6000, DS8000 port, XIV port	The average number of send operations and receive operations per second. Specify a threshold to receive alerts about storage system ports that might have too high a workload.
Zero Buffer Credit Percentage	SAN Volume Controller port	<p>The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>Warning boundary: greater than 25%</li> <li>Critical boundary: greater than 66%</li> </ul>
Zero Buffer Credit Rate	Brocade switch port	<p>The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>Warning boundary: greater than 75,000,000 milliseconds</li> <li>Critical boundary: greater than 198,000,000 milliseconds</li> </ul>

Table 140. Port thresholds (continued)

Threshold (Metric)	Resource Type	Description
Zero Buffer Credit Timer	SAN Volume Controller port	<p>The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.</p> <p>This threshold has the following default boundaries:</p> <ul style="list-style-type: none"> <li>• Warning boundary: greater than 75,000,000 milliseconds</li> <li>• Critical boundary: greater than 198,000,000 milliseconds</li> </ul>
Zero Receive Buffer Credit Percentage	DS8000 port	<p>The amount of time, as a percentage, that the port on DS8000 had depleted its receive buffer credits. That is, the percentage of time that the port on DS8000 had no credit to provide to the sending port.</p>
Zero Send Buffer Credit Percentage	DS8000 port	<p>The amount of time, as a percentage, that the port on DS8000 had depleted its send buffer credits. That is, the percentage of time that the receiving port had no credit to provide to the port on DS8000.</p>

## Cinder driver for Tivoli Storage Productivity Center

The Cinder driver for Tivoli Storage Productivity Center enables cloud storage that is powered by OpenStack to use a Tivoli Storage Productivity Center installation for block storage provisioning.

Tivoli Storage Productivity Center provides block storage provisioning capabilities that the storage administrator can employ to define the properties and characteristics of storage volumes within a particular service class. For example, a block storage service class can define RAID levels, tiers of storage, and various other storage characteristics. For more on service classes, see “Configuring Tivoli Storage Productivity Center for provisioning” on page 245.

The Cinder driver enables an OpenStack cloud to use these defined service classes, which are available in your Cinder environment as Cinder volume types. For each block storage service class, a Cinder volume type is created. Your cloud users can create and use volumes of this type.

Download the Tivoli Storage Productivity Center Cinder driver and installation instructions from <http://www.ibm.com/developerworks/servicemanagement/sm/tpc/index.html>.

## Planning to deploy the Cinder driver

Before using the Cinder driver for Tivoli Storage Productivity Center you need to create volume types and service classes for the service to use. You also need to decide the connection protocol your storage will be using (FC or iSCSI).

### Volume types within OpenStack

New volumes are created within OpenStack by specifying the type of volume to be created. The volume types are high-level descriptions that identify the class of storage a volume can provide. For instance, a cloud user might want to create a new volume with Gold class storage. The user can employ the OpenStack Horizon GUI or Cinder CLI to provision a new volume. The user can then specify that the type of volume they want to create is of type Gold. The OpenStack Cinder scheduler then allocates from those volume drivers that can provision a volume of type Gold.

The Cinder driver works with this information through the medium of service classes. A service class is used to identify classes of storage within the storage infrastructure that can be configured for such services as thin provisioning, virtualized volumes, and varying RAID levels. Service classes can be configured to use storage across multiple devices and device types that still meet the constraints that are defined within the service class. The Cinder driver can discover the service classes that are configured within Tivoli Storage Productivity Center and create corresponding volume types within the OpenStack environment. The user can then provision a volume from the service classes that are already defined within Tivoli Storage Productivity Center.

### Creating a service class and volume type

To provide a new class or type of volume, the storage administrator creates a service class within Tivoli Storage Productivity Center. The definition includes the characteristics of the service class and the storage pools that are available to the service class. When the service class is created, the Cinder driver detects the new service class and automatically creates a corresponding Volume Type entry within the OpenStack infrastructure.

### Specifying connection protocols for storage (FC and iSCSI)

When the Cinder driver is requested to attach a volume to a Nova node, the Cinder driver submits required access information to the Tivoli Storage Productivity Center service. The service then assigns the volume to all available host ports associated with the Nova node.

When the Tivoli Storage Productivity Center service assigns the volume to the ports, the service returns all available access information back to the Cinder driver. The Cinder driver chooses the optimal access to give to the Nova node to initialize

the connection. By default, the Cinder driver attempts to choose the Fibre Channel method of access. The Cinder driver determines whether the Tivoli Storage Productivity Center server can enable Fibre Channel access and if the Nova node can use Fibre Channel access. If both are available, Fibre Channel connection information is returned. If Fibre Channel access is not available, the Cinder driver checks whether the Tivoli Storage Productivity Center service is able to assign iSCSI access to the volume. If iSCSI access is available to the Nova node, the Cinder driver returns the iSCSI access information to the Nova node.

An OpenStack cloud administrator can change which access protocol the Cinder driver attempts to assign first. The setting is the `tpc_preferred_protocol` parameter in the `cinder.conf` driver configuration section. By changing the preferred protocol to iSCSI, the Cinder driver first attempts to provide iSCSI access. If unsuccessful, the Cinder driver falls back to Fibre Channel access. If the preferred protocol cannot be met but an alternative protocol is available, then the alternative protocol is chosen.

## Enabling an OpenStack cloud for Tivoli Storage Productivity Center

Enabling your OpenStack cloud to leverage your Tivoli Storage Productivity Center installation is a two step process. Configuration scripts for enabling Tivoli Storage Productivity Center functionality and policies for the Cinder driver are available in the `tpc-scripts` directory.

Download this zip file and extract its contents to a location on the Tivoli Storage Productivity Center server. Open a command prompt to the `tpc-scripts` directory and run the setup script for your platform.

**Note:** You must restart the Tivoli Storage Productivity Center Device, Data, and Web servers to complete the setup process.

A natively installable package containing the Cinder driver is available in the `cinder-packages` directory. Copy the appropriate `.rpm` or `.deb` file to your Cinder node and install it using the operating system's native packaging tool.

For example, if the Cinder node is running Red Hat, copy the `openstack-tpcdriver-1.0.0-1.x86_64.rhel.rpm` file and install the file using the command `rpm -ivh openstack-tpcdriver-1.0.0-1.x86_64.rhel.rpm`.

Upon installation of the `.rpm` or `.deb` file, you can run the `tpcdriver-config.sh` utility to help determine appropriate Cinder settings for your installation.

Copy the appropriate settings into your `cinder.conf` file and restart the Cinder volume service to enable the Cinder driver.

Once the Cinder driver is installed and configured, you can verify the installation and configuration by navigating to the Volumes page of the Horizon Admin panel to check that the service classes configured within Tivoli Storage Productivity Center are now available as volume types.

## Setting up the Cinder driver environment

To use the Cinder driver, you need to configure role-to-group mappings for an authorized Cinder user, add storage system mappings, and define capacity pools and service classes.

## Configuring volume creation and attachment

After you install the Tivoli Storage Productivity Center, but before you configure the Cinder node, use the Tivoli Storage Productivity Center web-based GUI to set up volume creation and attachment.

1. Log in to the web-based GUI.
2. Configure role-to-group mappings for an authorized OpenStack Cinder user in a selected group to connect Tivoli Storage Productivity Center.
  - a. In the menu bar, go to **Settings > User Management**.
  - b. On the **User Management** window, click **Add Group**.
  - c. On the **Add Group** window, enter the name of the group to which the user belongs, and click **Search**.
  - d. Select the group name and assign the **External Application** role to this group.
  - e. Click **OK**.
3. Add storage subsystems to manage to Tivoli Storage Productivity Center. For more information, see “Adding storage systems” on page 132.
  - a. In the menu bar, go to **Storage > Storage Systems**.
  - b. Click **Add Storage System**.
  - c. Click the **Storage System** icon.
  - d. On the **Discover** page, enter the information, and click **Next**.
  - e. On the **Configure** page, enter the information, and click **Configure**.
4. Define capacity pools for the storage subsystems to be used by the service classes. For more information, see “Capacity pools” on page 74.
  - a. In the menu bar, go to **Storage > Pools**.
  - b. Select one or more pools of a storage system and select **Action > Add to Capacity Pool**.
  - c. Click **Create Capacity Pool**.
  - d. In the **Create Capacity Pool** page, specify a name for the capacity pool.
  - e. Click **OK**.
5. Define the service classes to be used by the Cinder driver. For more information, see “Service classes” on page 70.
  - a. In the menu bar, go to **Advanced Analytics > Cloud Configuration**, and click **Work With Service Classes**.
  - b. Click **Create Service Class**.
  - c. In the **Select Type** page, click on the **Block** icon to create a block storage service class.
  - d. In the **Define Properties** page, specify a name for the service class.
  - e. Specify the requirement for provisioning block-storage of this service class, and click **Next**.
  - f. In the **Capacity Pool** page, select **Only allow provisioning from selected capacity pools**, select the capacity pools, and click **Next**.
  - g. In the **Authorize Users** page, disable **Approval required**. Select **Allow administrators and specified users to provision**.
  - h. Select the user name
  - i. Click **Add User**, and then click **Finish**.
6. Install the Tivoli Storage Productivity Center OpenStack patch and the Cinder driver. See “Installing the Cinder driver” on page 1132.

7. Use the OpenStack Horizon GUI to verify that the service classes are listed as volume types on Cinder nodes.
  - a. Log in to the OpenStack Horizon GUI as a user with admin roles.
  - b. Select **Admin tab > System Panel > Volumes**.
  - c. Verify that the service classes created on Tivoli Storage Productivity Center are listed under Volume Types.
  - OR
  - a. Log in to the OpenStack Cinder node.
  - b. Use the CLI cinder command to list volume type: # cinder type-list.
  - c. Verify that the service classes you created on Tivoli Storage Productivity Center are listed.
  - d. Ensure that the correct settings for Tivoli Storage Productivity Center are added to the Cinder configuration file /etc/cinder/cinder.conf. Check that the Cinder volume service was restarted.
8. Optionally, you can configure volume snapshot creation. To create a volume snapshot with the Cinder driver, add the storage systems to Tivoli Storage Productivity Center for Replication.
  - a. In the menu bar in the web-based GUI, go to **Storage > Storage Systems**. Log in to Tivoli Storage Productivity Center web-based GUI.
  - b. On the Storage Systems page, click **Add Storage System**.
  - c. Click the icon for the type of storage system that you want to add.
  - d. Complete the connection information for the storage system and schedule the collection of data.
  - e. Ensure that the same storage systems are also added to Tivoli Storage Productivity Center
9. iSCSI volume attachment

To use iSCSI connections, the storage system must be configured for iSCSI access. You must assign IP addresses to specific Ethernet ports on the storage subsystem itself. For each storage subsystem, consult the product documentation on how to configure Ethernet ports for iSCSI access.

After the storage system is configured for iSCSI access, the Tivoli Storage Productivity Center server populates information about the iSCSI access with the next probe of the storage subsystem. You can manually trigger this action from the Tivoli Storage Productivity Center GUI by running the Start Probe command under the **Data Collection** menu for storage subsystems.

When you assign a volume to a server with iSCSI host ports, the Tivoli Storage Productivity Center generates and assigns a pseudo-random 16 character CHAP secret to any host definitions it creates on the storage subsystem. If the host port already belongs to a host definition on the storage system, the server does not carry out any CHAP secret modification.

**Note:** In Tivoli Storage Productivity Center version 5.2.1, iSCSI ports appear as Fibre Channel ports within the Tivoli Storage Productivity Center GUI. The WWPN of the iSCSI port is the iSCSI name of the port itself.

In Tivoli Storage Productivity Center 5.2.2, iSCSI ports appear as Fibre Channel ports within the Tivoli Storage Productivity Center GUI. The WWPN of the iSCSI port is the iSCSI name of the port itself.

## Requirements for installing the Cinder driver

To offer cloud services for provisioning block files, you must ensure that you have the correct versions of Tivoli Storage Productivity Center and OpenStack.



## Requirements for Tivoli Storage Productivity Center

You need to install Tivoli Storage Productivity Center version 5.2.2 or later with Virtual Storage Center license. You also need:

- A local operating system, LDAP, or Windows Domain group to add to Tivoli Storage Productivity Center
- A user added as a member of this group for use by the Cinder driver to connect to Tivoli Storage Productivity Center
- A supported version of DB2 client or DB2 server
- Service classes available to the OpenStack Cinder node(s) as volume types
- Storage subsystems available for OpenStack volume creation and attachment support
- Storage subsystems configured in Tivoli Storage Productivity Center for Replication for OpenStack volume snapshot support

### Storage platform support

The following storage platforms are supported:

- DS8000
- IBM SAN Volume Controller 6.1 or later
- Storwize V7000 6.1 or later
- XIV 11.0 or later

## Requirements for OpenStack

You will need:

- OpenStack (Havana release)
- The OpenStack Cinder node must run on one of the following operating systems:
  - Red Hat Enterprise Linux 6.4 or higher
  - Ubuntu 12.04
- The python-cinder (included in python-cinderclient package) and python-rdfliib library packages installed on the Cinder node:
  - Red Hat Enterprise Linux 6.4 or higher.

The Red Hat python-rdfliib library is available in the Extra Packages for Enterprise Linux (EPEL) repository. Visit the [Fedora wiki](#) for information on installing the EPEL repository on your system.

```
# yum install python-cinderclient
# yum install python-rdfliib
```

- Ubuntu 12.04

```
# sudo apt-get install python-cinderclient
# sudo apt-get install python-pip
# sudo pip install rdfliib==3.2.1
```

**Note:** The python-rdfliib debian package available in the Ubuntu precise archives contains rdfliib 2.4.2. The Tivoli Storage Productivity Center Cinder driver requires version 3.2.1 or higher of the rdfliib library. Use the pip command above in order to install the correct version of the rdfliib library.

## Cinder driver package contents

To enable an OpenStack cloud to leverage a Tivoli Storage Productivity Center installation, run the installation script and use a configuration utility to determine appropriate Cinder settings for your installation.

The `tpc-scripts` directory contains the platform-specific Tivoli Storage Productivity Center Server OpenStack scripts to enable or disable the Tivoli Storage Productivity Center functionality and policies for the Cinder driver.

- AIX Scripts: `setup_aix.sh` and `unsetup_aix.sh`
- Linux Scripts: `setup_linux.sh` and `unsetup_linux.sh`
- Windows Scripts: `setup_windows.bat` and `unsetup_windows.bat`

The `setup` script configures a Tivoli Storage Productivity Center server with the files required to enable the supported OpenStack functionality on the Tivoli Storage Productivity Center server. For more information about installing the Cinder driver, see “Requirements for installing the Cinder driver” on page 1130 and “Installing the Cinder driver.”

The `unsetup` script reconfigures a Tivoli Storage Productivity Center server with the files required to disable supported OpenStack functionality on the a Tivoli Storage Productivity Center server. For more information about uninstalling the Cinder driver, see “Uninstalling the Cinder driver” on page 1134.

When Tivoli Storage Productivity Center Cinder driver is installed, you are ready to enable OpenStack and Tivoli Storage Productivity Center to work together. The `cinder-packages` directory contains the platform specific Cinder driver installer packages . Both the Tivoli Storage Productivity Center server and the OpenStack Cinder node need to be patched.

- Red Hat Enterprise Linux 6.4 or higher: `openstack-tpcdriver-1.0.1-1.x86_64.rhel.rpm`
  - Ubuntu 12.04: `openstack-tpcdriver_1.0.1-1_all.deb`
1. Log on to Tivoli Storage Productivity Center server as root for Linux/AIX or as Administrator for Windows.
  2. Change directory to `tpc-scripts`.
  3. Run the `setup` script, providing the path to the Tivoli Storage Productivity Center installation location. For example:

Option	Description
AIX	<code># ./setup_aix.sh /opt/IBM/TPC</code>
Linux	<code># ./setup_linux.sh /opt/IBM/TPC</code>
Windows	<code>&gt; setup_windows.bat "C:\Program Files\IBM\TPC"</code>

4. The installer package installs the OpenStack `tpc.py` driver in the proper location. It also installs the `tpcdriver-config.sh` script which generates the additional Cinder driver-related entries needed for the Cinder configuration file (`/etc/cinder/cinder.conf`).

## Installing the Cinder driver

To install the Cinder driver, run the installation script and use a configuration utility to determine appropriate Cinder settings for your installation.

The Tivoli Storage Productivity Center server and all OpenStack Cinder nodes need to be patched to enable Tivoli Storage Productivity Center OpenStack functions and policies for the Cinder driver. Review “Requirements for installing the Cinder driver” on page 1130 before you proceed with the installation.

1. Run the Tivoli Storage Productivity Center OpenStack patch installation script for AIX, Linux, or Windows. On Tivoli Storage Productivity Center:
  - a. Log on as root for Linux/AIX or **Administrator** for Windows.
  - b. Change directory to tpc-scripts.
  - c. Run the setup script, specifying the path to the Tivoli Storage Productivity Center server installation location:
    - AIX# ./setup\_aix.sh /opt/IBM/TPC
    - Linux# ./setup\_linux.sh /opt/IBM/TPC
    - Windows: setup\_windows.bat "C:\Program Files\IBM\TPC"
  - d. When prompted if you want to proceed with the update, enter Y or y.

**Note:** You must restart the Tivoli Storage Productivity Center Device, Data, and Web servers to complete the setup process.

2. The Tivoli Storage Productivity Center OpenStack patch installation script attempts to determine the DB2 settings. If the script cannot find a setting, you are prompted to enter the value. Descriptions of the settings are in the following table:

Option	Description
<b>Hostname of the DB2 server</b>	The host name of the Tivoli Storage Productivity Center database server.
<b>Port of the DB2 server</b>	The port that the Tivoli Storage Productivity Center database server is running on.
<b>DB2 profile user (AIX/Linux only)</b>	The user name of the DB2 user on the server this installation script is being run on. This setting is only relevant for AIX and Linux.
<b>DB2 user</b>	The user name of the DB2 user with permissions to access the Tivoli Storage Productivity Center database.
<b>DB2 password</b>	Password for the DB2 user.
<b>Database name</b>	The name of the Tivoli Storage Productivity Center database.
<b>Schema name</b>	The schema name of the Tivoli Storage Productivity Center database.

3. Run the Cinder driver installation script for Red Hat and Ubuntu installations. On the OpenStack Cinder nodes:
  - a. Ensure that the current user has administrative privileges - **root** for Red Hat and Ubuntu.
  - b. Change directory to cinder-packages.
  - c. Install the package for your OpenStack distribution:
    - Red Hat Enterprise Linux 6.4: # rpm -ivh openstack-tpcdriver-1.0.0-1.x86\_64.rhel.rpm
    - Ubuntu 12.04: # dpkg -i openstack-tpcdriver\_1.0.0-1.x86\_64.deb

**Note:** If you installed any packages manually (for example, with pip or from source), you can force the installation by overriding dependency checks. On Red Hat systems, the --nodeps flag is passed to the rpm command. On Ubuntu systems, the --force-depends flag is passed to the dpkg command.

4. Run the `tpcdriver-config.sh` configuration script. You are prompted to provide Tivoli Storage Productivity Center server information as the script runs: `# /usr/local/bin/tpcdriver-config.sh` Refer to “Configuring the Cinder driver” for more details on the Cinder configuration settings specific to Tivoli Storage Productivity Center.
5. Copy the Cinder driver settings that are generated by the `tpcdriver-config.sh` script to the [DEFAULT] section of the Cinder configuration file, `/etc/cinder/cinder.conf`.
6. Restart the Cinder volume service.

## Uninstalling the Cinder driver

To uninstall the Cinder driver, run uninstallation scripts for AIX, Linux, and Windows installations.

Use Tivoli Storage Productivity Center OpenStack patch uninstallation scripts for AIX, Linux, and Windows to uninstall an instance of the Cinder driver.

1. Log on to Tivoli Storage Productivity Center with administrative privileges - **root** for Linux/AIX or **Administrator** for Windows.
2. Change directory to `tpc-scripts`.
3. Run the `unsetup` script, specifying the path to the Tivoli Storage Productivity Center server installation location:
  - AIX
 

```
# ./unsetup_aix.sh /opt/IBM/TPC
```
  - Linux
 

```
# ./unsetup_linux.sh /opt/IBM/TPC
```
  - Windows
 

```
unsetup_windows.bat "C:\Program Files\IBM\TPC"
```
4. When prompted if you want to proceed with the update, enter **Y** or **y**.
5. You must restart the Tivoli Storage Productivity Center Device, Data, and Web servers to complete the `unsetup` process.
6. On the OpenStack Cinder node(s), log on as **root** for Red Hat and Ubuntu.
7. Configure the Cinder volume service to no longer use the Tivoli Storage Productivity Center Cinder volume driver by editing the Cinder configuration file `/etc/cinder/cinder.conf` and removing all references to the Tivoli Storage Productivity Center Cinder volume driver. Parameters for the driver are prefixed with the string `tpc_`.
8. Restart the Cinder volume service.
9. After the Cinder volume services are no longer configured to use the Tivoli Storage Productivity Center driver, the native package can be removed from the Cinder volume node. For example:
  - Red Hat
 

```
# rpm -e openstack-tpcdriver
```
  - Ubuntu
 

```
# dpkg -r openstack-tpcdriver
```

## Configuring the Cinder driver

To make Tivoli Storage Productivity Center volumes and service classes available to Cinder services, you must configure the Cinder driver.

On Tivoli Storage Productivity Center, you must plan user and service class details:

- Identify which Tivoli Storage Productivity Center user the Cinder driver can use to log on. The user's group must be added to **User Management** as **External Application**.
- Identify the preferred service class. This class is used when no volume type is specified in the Cinder configuration file. All block service classes available to the user are available to the Cinder driver as volume types.

The `tpcdriver-config.sh` script installed on the Cinder node as part of the rpm or deb file helps to configure the Cinder volume service to use the Cinder driver. Configure the Cinder driver by editing the Cinder configuration file (`/etc/cinder/cinder.conf`).

Edit the following `volume_` settings in the Cinder configuration file:

- `volume_driver`: The name of the Cinder driver to use for volume creation. By default, name is set as `cinder.volume.drivers.lvm.LVMISCSIDriver` in the Cinder configuration file. To use the Cinder driver, change the setting to `cinder.volume.drivers.tpc.TPCDriver`.
- `volume_name_template`: The template that is used by Tivoli Storage Productivity Center to name OpenStack volumes. By default, template is set as `volume-%s` in the Cinder configuration file. You can change this value to one more easily identifiable as an OpenStack volume (for example, `openstackvol-%s`) when you are viewing or searching volumes. The template is a global setting that affects volume names for all Cinder drivers.

You must add the following `tpc_` settings to the Cinder configuration file:

- `tpc_ip`: The IP address of the Tivoli Storage Productivity Center server.
- `tpc_login`: The user name that the Cinder driver will log on with.
- `tpc_password`: The password that the Cinder driver will log on with.
- `tpc_service_class`: The default Tivoli Storage Productivity Center service class that the Cinder driver uses if no volume type is specified.
- `tpc_preferred_protocol`: The preferred protocol if more than one is available. The preferred protocol can be set to `iSCSI` or `FC`: these values are case-sensitive. This flag is optional - if not specified, the default is `FC`.

After you edit the Cinder configuration file, restart Cinder services with the command: `service service-name restart`. For Ubuntu, the `service-name` parameter is `cinder-volume`. For RedHat and SLES, the `service-name` parameter is `openstack-cinder-volume`.

Use the Horizon GUI to create, attach, and delete volumes.

## Learn more about the Cinder driver

You can learn more about the Cinder driver from Service Management Connect resources such as developer wikis, updates, and Beta packages.

Tivoli Storage Productivity Center Service Management Connect contains the OpenStack package and links to

- Tivoli Storage Productivity Center developerWorks wiki
- Tivoli Storage Productivity Center forums
- Updates to this package
- Tivoli Storage Productivity Center Beta material

To learn more, go to <http://www.ibm.com/developerworks/servicemanagement/sm/tpc/index.html>

The IBM Knowledge Center contains information on how to use and configure your Tivoli Storage Productivity Center installation and other IBM products. To learn more, go to [http://www.ibm.com/support/knowledgecenter/SSNE44\\_5.2.4/](http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.4/).

The OpenStack website contains information on installing, configuring, and managing an OpenStack cloud. To learn more, go to <http://www.openstack.org>.

## Limitations and known issues for the Cinder driver

The Cinder driver is subject to the following limitations and known issues.

### Decommissioning a service class and volume type

To remove a volume type from OpenStack, you must first remove the service class that is associated with the volume type in Tivoli Storage Productivity Center. After the service class is removed, the OpenStack Administrator must then remove the corresponding volume type in OpenStack.

For example, you can remove a service class (Gold) from Tivoli Storage Productivity Center, and then remove the type of volume that is associated with the Gold service class. If a user requests the type of volume that is associated with the Gold service class, the request is rejected.

If you remove the volume type from OpenStack but do not remove the associated service class, the volume type is re-created the next time the Cinder driver checks for available service classes.

### iSCSI support with XIV Storage System

- For an XIV system with multiple Internet Small Computer System Interface (iSCSI) ports, only one port is populated by the Cinder driver.
- For the XIV system, Tivoli Storage Productivity Center is unable to use Challenge-Handshake Authentication Protocol (CHAP) secrets for iSCSI host definitions. CHAP information is not available through the CLI for the XiV system. iSCSI attachments that are made through the Cinder driver are assigned without CHAP secrets defined.

### Mixed iSCSI and Fibre Channel Volume Assignments

Tivoli Storage Productivity Center attempts to assign all ports that belong to a Nova node. When the unmanaged server representing the Nova node contains both Fibre and iSCSI ports, the volume is available to the Fibre Channel and iSCSI ports on the Nova node. In some cases, when you detach a volume from a host, the ports for the access method in use by the Nova node are removed but the access for the alternative connectivity method remains. This access can cause errors to delete the volume from the storage system. In this case, the volume is removed from the Cinder database but it requires manual steps to remove the storage volume from the storage system.

## SAN Volume Controller / Storwize Snapshot and Volume Copies

In some cases, a snapshot session is removed from the Tivoli Storage Productivity Center for Replication configuration while a background copy is still in progress. The snapshot or volume copy is not affected by this configuration removal, but a delete snapshot or volume action fails if the FlashCopy relationship remains configured on the hardware. To resolve this problem, remove the FlashCopy relationship from the storage virtualizer and manually delete the volume.

### Unable to Attach Storage Volumes

In some circumstances, attempts to attach a storage volume to a compute instance fails even though the volume is assigned to the Nova node's host connection information. To resolve this problem, verify that the Nova node can connect to the storage system through the Ethernet LAN or Fibre SAN. If an XIV Storage System is added (regardless of whether the volume is on an XIV Storage System), restart the Tivoli Storage Productivity Center Device server. Use the startup scripts in the scripts directory where Tivoli Storage Productivity Center is installed.

---

## Opening Tivoli Storage Productivity Center on Windows operating systems

You can open Tivoli Storage Productivity Center CLIs and GUIs and administer Tivoli Storage Productivity Center on Windows operating systems.

### Opening Tivoli Storage Productivity Center CLIs and GUIs

To manage and monitor storage resources, open Tivoli Storage Productivity Center CLIs and GUIs.

You can open the following GUIs and CLIs:

- “Opening Tivoli Storage Productivity Center Web-based GUI”
- “Opening Tivoli Storage Productivity Center Stand-alone GUI” on page 1138
- “Opening Tivoli Storage Productivity Center Replication Manager GUI” on page 1138
- “Opening DB2 Command Window” on page 1138
- “Opening IBM Data Studio Administration client on Windows operating systems” on page 1139
- “Opening IBM Tivoli Monitoring Services” on page 1139

### Opening Tivoli Storage Productivity Center Web-based GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none"><li>1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b>.</li><li>2. On the Start page, right-click, and then click the <b>All apps</b> taskbar button.</li></ol>
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start &gt; All Programs</b> .

2. Click **Tivoli Storage Productivity Center > TPC Web-based GUI**.

## Opening Tivoli Storage Productivity Center Stand-alone GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none"><li>1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b>.</li><li>2. On the Start page, right-click, and then click the <b>All apps</b> taskbar button.</li></ol>
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start &gt; All Programs</b> .

2. Click **Tivoli Storage Productivity Center > TPC Stand-alone GUI**.

## Opening Tivoli Storage Productivity Center Replication Manager GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none"><li>1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b>.</li><li>2. On the Start page, right-click, and then click the <b>All apps</b> taskbar button.</li></ol>
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start &gt; All Programs</b> .

2. Click **Tivoli Storage Productivity Center > TPC Replication Manager GUI**.

## Opening DB2 Command Window

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none"><li>1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b>.</li><li>2. On the Start page, right-click, and then click the <b>All apps</b> taskbar button.</li></ol>
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start &gt; All Programs</b> .

2. Click **IBM DB2 > Command Line Tools > Command Window**.



## Opening IBM Data Studio Administration client on Windows operating systems

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none"><li>1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b>.</li><li>2. On the Start page, right-click, and then click the <b>All apps</b> taskbar button.</li></ol>
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start &gt; All Programs</b> .

2. Click **IBM Data Studio > Data Studio Administration Client**.

## Opening IBM Tivoli Monitoring Services

1. Choose one of these options:

Option	Description
Windows Server 2012	<ol style="list-style-type: none"><li>1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b>.</li><li>2. On the Start page, right-click, and then click the <b>All apps</b> taskbar button.</li></ol>
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start &gt; All Programs</b> .

2. Click **IBM Tivoli Monitoring > IBM Tivoli Monitoring Services**.

## Accessing administration tools

To manage and maintain Tivoli Storage Productivity Center, access the Windows system administration tools.

To complete tasks in Tivoli Storage Productivity Center, you must open the following administration and maintenance facilities:

- “Accessing the Control Panel”
- “Accessing Administrative Tools” on page 1140
- “Accessing Windows Services” on page 1140
- “Accessing Computer Management” on page 1140
- “Accessing Programs and Program Features” on page 1140

## Accessing the Control Panel

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b> .

Option	Description
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start</b> .

2. Click **Control Panel**

## Accessing Administrative Tools

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b> .
Windows 7, Windows 2008	Click <b>Start</b> .
Windows XP, Windows Vista	Click <b>Start &gt; Control Panel</b> .

2. Click **Administrative Tools**

## Accessing Windows Services

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b> .
Windows 7, Windows 2008	Click <b>Start</b> .
Windows XP, Windows Vista	Click <b>Start &gt; Control Panel</b> .

2. Click **Administrative Tools > Services**

## Accessing Computer Management

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b> .
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start</b> .

2. Click **Control Panel > Computer Management**

## Accessing Programs and Program Features

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b> .

Option	Description
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start</b> .

2. Click **Control Panel > Programs and Program Features**.

### Accessing Window Run

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the <b>Server Manager</b> taskbar button, and then click <b>Start</b> .
Windows 7, Windows 2008, Windows Vista, Windows XP	Click <b>Start</b> .

2. Click **Run**.




---

## Appendix B. Accessibility features for Tivoli Storage Productivity Center

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

### Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Productivity Center:

- Keyboard-only operation in the stand-alone GUI and the web-based GUI.  
Restriction: In the stand-alone GUI, you must use the mouse to navigate the Topology Viewer and report graphs.
- Interfaces that are commonly used by screen readers.
- An information center that includes the following accessibility features:
  - The information center is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
  - All documentation for Tivoli Storage Productivity Center is available in Adobe Portable Document Format (PDF) by using the Adobe Acrobat Reader. To access PDFs, go to  <ftp://public.dhe.ibm.com/software/tivoli/tpc/v524/>.
  - All images in the information center are provided with alternative text, so that visually impaired users can understand the contents of the images.

### Keyboard navigation in the stand-alone GUI

The stand-alone GUI uses standard Microsoft Windows navigation keys. However, you must use the mouse to navigate the Topology Viewer and report graphs. Standard operating system keystrokes are used for standard operating system operations.

### Keyboard navigation in the web-based GUI

Most of the features of the web-based GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the web-based GUI:

#### For navigating the web-based GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

#### For actions menus:

- To navigate to the grid header, press Tab.

- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

**For filters:**

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

**For text fields:**

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

**For tables or lists:**

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

**Restriction:** For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

## **IBM and accessibility**

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at  <http://www.ibm.com/able/>.

---

## Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

### Accessibility features

The following list includes the major accessibility features in Tivoli Storage Productivity Center for Replication:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices
- An information center that includes the following accessibility features:
  - The information center is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
  - All documentation for Tivoli Storage Productivity Center for Replication is available in Adobe Portable Document Format (PDF) by using the Adobe Acrobat Reader. To access the PDFs, go to <ftp://public.dhe.ibm.com/software/tivoli/tpc/v524/>.
  - All images in the information center are provided with alternative text so that visually impaired users can understand the contents of the images.

### Keyboard navigation

Use the following key combinations to navigate the interface by keyboard:

- To go directly to the topic pane, press Alt+K, and then press Tab.
- In the topic pane, to go to the next link, press Tab.
- To go directly to the Search Results view, press Alt+R, and then press the Enter or Up-Arrow key to enter the view.
- To go directly to the Navigation (Table of Contents) view, press Alt+C, and then press the Enter or Up-Arrow key to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left-Arrow keys.
- To move to the next topic node, press the Down-Arrow or Tab key.
- To move to the previous topic node, press the Up-Arrow key or Shift+Tab.
- To go to the next link, button, or topic node from inside on of the views, press Tab.
- To scroll all the way up or down in a pane, press Home or End.
- To go back, press Alt+Left Arrow; to go forward, press Alt+Right Arrow.
- To go to the next pane, press F6.
- To move to the previous pane, press Shift+F6.
- To print the active pane, press Ctrl+P.

## Related accessibility information

The following list contains hints and tips that can help you more fully use the graphical user interface:

**Drop-down lists are positioned directly over or before the radio button that activates it.**

If you use a screen reader, you should be aware that there are radio buttons to activate drop-down lists for several GUI pages. The way to activate the drop-down list is by selecting the associated radio button. The drop-down list is positioned directly over or before the radio button that activates it. When you use a screen reader that processes the fields and controls of a page sequentially, you might select the radio button, but not know that the associated drop-down list has been activated. The screen reader processes inactive drop-down lists first, and then processes the next radio button. The drop-down list is activated if you select the radio button.

On the following pages, keep in mind that radio buttons activate a drop-down list:

- Administration
- ESS/DS Paths
- Sessions
- Session Details
- Storage Systems

**Tables are best understood by reviewing the surrounding text and the table row and column number of the table.**

On some graphical user pages, tables use the header or row ID attributes when reading a single cell. The screen reader reads the table row and column number, along with cell data. Therefore, you can infer the column header and row ID.

**Experiment with and fine-tune the way your screen reader pronounces some of the product abbreviations.**

Your screen reader might pronounce abbreviations as if they were words. For example, the common abbreviation for Enterprise Storage Server is ESS. Your screen reader might read ESS as the word "ess". With some screen readers you can hear alternate pronunciations. If you frequently use the software you might prefer to fine-tune such associations in your settings. When an association is created, the screen reader can recognize the abbreviation as a word. If you can add dictionary words with your screen reader, replace the capitalized character sequence with the sequence E space S space S.

Typically, this abbreviation is used in the combination form of ESS/DS. This term refers to TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, or System Storage DS8000.

**Some decorative artifacts might persist if the cascading style sheet is disabled.**

Enable cascading style sheets when possible; otherwise, some decorative elements might persist in the web browser GUI. These artifacts do not affect performance. If they become too distracting, consider using the command-line interface instead.

**For efficiency, confirmation dialogs place initial focus on the Yes button.**



When a confirmation dialog box is displayed, focus is given to the **Yes** button. Therefore, the screen reader reads “Yes” but does not read the confirmation text. The software processes the information in this way when you do the following types of tasks:

- Perform an action on a session
- Remove a connection to a storage system
- Click the **About link**
- Create a high-availability connection

To read the confirmation text before clicking the **Yes**, **No**, or **OK** button, view the previous heading before the button.

**Dojo components are not read by all screen readers.**

The Job Access for Windows and Speech (JAWS) screen reader does not read some Dojo components on Windows Internet Explorer. Use the command-line interface instead of the GUI with JAWS on Windows Internet Explorer.

**Firefox is the preferred browser for use with a screen reader.**

Use Firefox as the screen reader because other browsers might not fully expose assistive technology content to the screen reader.

## **IBM and accessibility**

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at  <http://www.ibm.com/able>.



---

## Notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan, Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation  
2Z4A/101  
11400 Burnet Road  
Austin, TX 78758  
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

**COPYRIGHT LICENSE:** This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these

programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

---

## Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and

IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

---

## Glossary

This glossary provides terms and definitions for the IBM Tivoli Storage Productivity Center software and products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website ([opens in new window](#)).

---

### A

#### active management server

A management server from which the storage environment can be monitored and managed. The active management server replicates its database to the standby server.

---

### B

#### business continuity

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

---

### C

**CIM** See Common Information Model.

#### CIM agent

The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

**CKD** See count key data.

#### cluster

1. In Storwize<sup>®</sup> V7000, a pair of nodes that provides a single configuration and service interface.

2. A loosely coupled collection of independent systems (or nodes) organized into a network for the purpose of sharing resources and communicating with each other.
3. A collection of complete systems that work together to provide a single, unified computing capability.
4. In IBM System Storage DS8000, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

#### Common Information Model (CIM)

An implementation-neutral, object-oriented schema for describing network management or systems management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

#### community name

The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

#### consistency group

A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

#### copy set

The set of source volumes or target volumes involved in a FlashCopy operation.

#### count key data (CKD)

A data recording format that uses self-defining record formats in which each record on a volume is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

---

## D

### data collection

The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include Domain Name System (DNS) probes, web page analyzers, or database analyzers. See also discovery.

### data source

A storage resource or agent that provides data about a storage environment.

### discovery

The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also data collection.

---

## E

**ECKD** See extended count key data.

### enterprise repository

A component of the data server that records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by the data server using Java Database Connectivity (JDBC).

**event** An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

### extended count key data (ECKD)

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

---

## F

### FlashCopy

An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a volume.

**freeze** An operation in which a storage system blocks I/O from the host system to the affected volumes on the primary site. A freeze operation stops mirroring between

the primary and secondary volumes to ensure data consistency at the secondary site. See also thaw.

---

## G

**GB** See gigabyte.

**GiB** See gibibyte.

### gibibyte (GiB)

In digital information storage, a base-2 unit of measurement equal to 1,073,741,824 bytes (2 to the 30th power). See also gigabyte.

### gigabyte (GB)

For processor storage, real and virtual storage, and channel volume, two to the power of 30 or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes. See also gibibyte.

### Global Copy

A non-synchronous long-distance copy option for data migration and backup. See also Global Mirror, Metro Mirror, remote mirror and copy.

### Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

### Global Mirror

An optional capability of the remote mirror and copy feature that provides a two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also Global Copy, Metro Mirror, remote mirror and copy.

### grain size

The unit size for allocating space on thin-provisioned volumes, such as 32, 64, 128, and 256 kibibyte. The grain size is defined when a volume is created.

**GUID** See Globally Unique Identifier.

---

## H

### Hardware Management Console (HMC)

In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.



**heat map**

A color-coded data chart in which colors are used to differentiate values in a data set.

**HMC** See Hardware Management Console.

**host volume**

A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing.

**hypervisor**

Software or a physical device that enables multiple instances of operating systems to run simultaneously on the same hardware.

---

**I****in-band discovery**

The process of discovering information about the storage area network (SAN), including topology and attribute data, through the Fibre Channel data paths. See also out-of-band discovery.

---

**J****journal volume**

A volume that holds a consistent copy of data until a new consistent copy is formed. The journal volume restores the last consistent point during a recovery.

---

**K**

**KB** See kilobyte.

**KiB** See kibibyte.

**kibibyte (KiB)**

A base-2 unit of measurement that is equal to 1,024 bytes (2 to the 10th power). See also kilobyte.

**kilobyte (KB)**

For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes. See also kibibyte.

---

**L****launch-in-context**

An operation in which a user starts a secondary application from a primary application to perform a specific task. Using the parameters, navigation instructions, and user credentials that are supplied by the primary application, the secondary application opens to the specific place in which to complete the task.

**logical unit number (LUN)**

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

**LUN** See logical unit number.

---

**M****Management Information Base (MIB)**

In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

**management server**

A system that provides a central point of control for managing data replication.

**management server relationship**

A connection between two replication servers, where one server acts as the active server and replicates the data that is necessary for the standby server to take control of the replication environment.

**MB** See megabyte.

**mebibyte (MiB)**

A base-2 unit of measurement that is equal to 1,048,576 bytes (2 to the 20th power). See also megabyte.

**megabyte (MB)**

For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes. See also mebibyte.

**Metro Global Mirror**

A three-site, high availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an

intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

### **Metro Mirror**

A function of the remote mirror and copy feature that constantly updates a secondary copy of a volume to match changes made to a source volume. See also Global Copy, Global Mirror, remote mirror and copy.

**MiB** See mebibyte.

**MIB** See Management Information Base.

---

## **N**

### **namespace**

The scope within which a Common Information Model (CIM) schema applies.

### **native interface**

An interface that is specific to a system or subsystem.

---

## **O**

### **out-of-band discovery**

The process of discovering storage area network (SAN) information, including topology and device data, without using the Fibre Channel data paths. A common mechanism for out-of-band discovery is the use of SNMP Management Information Base (MIB) queries, which are invoked over an IP network. See also in-band discovery.

---

## **P**

### **ping**

1. A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network.
2. The command that sends an Internet Control Message Protocol (ICMP) echo-request packet to a gateway, router, or host with the expectation of receiving a reply.

**pool** A grouping of storage space that consists of volumes, logical unit numbers (LUNs),

or addresses that share a common set of administrative characteristics.

### **practice volume**

A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

### **primary site**

A physical or virtual site that is made up of hardware, network, and storage resources. Typically, production operations run at the primary site. Data can be replicated to a secondary site for disaster recovery and failover operations. See also secondary site.

### **primordial pool**

Storage capacity that is unallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

**probe** A data collection job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units.

---

## **R**

### **recovery point objective**

The maximum amount of data loss that can be tolerated during a service interruption.

### **remote mirror and copy**

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also Global Copy, Global Mirror, Metro Mirror.

### **resource**

In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, and storage systems.

### **role**

A job function that identifies the tasks that a user can perform and the resources to which a user has access. A user can be assigned one or more roles.

### **role pair**

The association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the

association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

---

## S

**scan** A data collection job that monitors storage usage and file statistics on the resources in an environment.

**secondary site**

A physical or virtual site that is made up of the hardware, network, and storage resources that support the recovery needs of the primary site. When a failure occurs at the primary site, operations can continue at the secondary site. See also primary site.

**session**

A collection of source and target volumes that are managed to create consistent copies of data. The type of data replication that is associated with the session determines the actions that can be conducted for the volumes.

**site awareness**

The association of a location with each storage system in a session. Site awareness ensures that a volume can only be selected for a session if it matches the location of the site. Site awareness helps to prevent both reversing a hardware relationship and selecting volumes at the wrong location.

**SMI-S** See Storage Management Initiative Specification.

**standby management server**

A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

**Storage Management Initiative Specification (SMI-S)**

A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

**storage pod**

A logical entity of components of a system consisting of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

**storage resource group**

A named collection of logically related resources that are monitored by Tivoli Storage Productivity Center. Monitored resources can include fabrics, switches, computers, storage systems, and other storage resource groups.

---

## T

**target volume**

A volume that receives data from a host volume or another intermediate volume.

**thaw**

An operation in which a storage system releases the block of the I/O from the host system to the affected volumes on the primary site. A thaw operation can occur after a freeze operation ends and consistency is formed on the secondary site. See also freeze.

**thin provisioning**

A mechanism that provides the ability to define logical volume sizes that are larger than the physical capacity installed on the system.

**track space-efficient volume (TSE volume)**

A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume to the target volume.

**TSE volume**

See track space-efficient volume.

---

## V

**virtualization**

The substitution of virtual resources for actual resources, where the virtual resources have the same functions and external interfaces as their counterparts, but differ in attributes, such as size, performance, and cost. Virtualization is commonly applied to physical hardware resources by combining multiple physical resources into shared pools from which users receive virtual resources.

**virtual storage area network (VSAN)**

A fabric within the storage area network (SAN).

**VSAN** See virtual storage area network.

---

**Z**

**zone** A logical grouping of switches, switch ports, and their attached devices in a fabric.

**zone alias**

A name that is given to a collection of one or more zone members to be managed together.

**zone set**

A group of zones that function together on a fabric.

---

# Index

## A

- about this document xiii
- access path host replacement 251
- access privileges
  - granting 480
  - modifying 481
  - removing 482
  - viewing 481
- accessibility features
  - for Tivoli Storage Productivity Center 1143
  - for Tivoli Storage Productivity Center for Replication 1145
- active management servers
  - about 308
  - reconnecting 315
  - reinstalling during active session 314
  - taking over 315
- activity deviation percentage
  - formula 278
- activity levels
  - formula 278
- activity limits
  - formula 278
- adding 227, 229, 234
  - copy sets to a FlashCopy session 414
  - copy sets to a Global Mirror session 422
  - copy sets to a Metro Global Mirror session 425
  - copy sets to a Metro Mirror session 419
  - copy sets to a Snapshot session 417
  - ESS/DS series logical paths 337
    - using a CSV file 337
  - host system connections 333
  - SNMP managers 316
  - storage connections
    - DS6000 326, 327
    - DS8000 324, 325
    - SAN Volume Controller 326
    - Storwize 326
    - TotalStorage Enterprise Server 326, 327
    - z/OS 327
- adding applications 227
- adding hypervisors
  - ESX and ESXi 136
  - vCenter Server system 136
- adding resources 127, 229
  - fabrics 133
  - hypervisors 136
  - servers with agents 142
  - servers without agents 139
  - storage systems 132
  - switches 133
  - vCenter Server system 137
- adding subdepartments 234
- administering
  - users 478
- agents
  - CIM 88
- alert
  - performance threshold 110
- alert conditions and violations 159
- alert log 182
- alert notification settings
  - modifying 176
- alerting 159
  - alert log 159, 182
  - alerting
    - conditions and violations 159
    - notifications 159
    - overview 80, 1035
- alerts
  - acknowledging in the web-based GUI 184
  - actions 1065
  - configuring for performance monitors 155
  - configuring for probes 149
  - notifications 1065
  - removing from the web-based GUI 184
  - triggering conditions 1036
- alias 86
- analysis period
  - performance data, change 281
- analyze tiering
  - analysis period, change 281
  - conditions, setting 97, 99
  - hypervisors 271
  - servers 271
  - storage pools 272
  - storage virtualizers 272
  - thresholds, modify 274
  - volume placement 97, 99
  - volumes 272
  - volumes, colocate 101
- analyze tiering wizard 95
- application
  - scenario 106
- application context 231
- applications
  - adding applications to departments for monitoring 227
  - creating applications for monitoring 226
  - monitoring 225
  - removing 231
  - removing applications for monitoring 231
  - resources 231
  - scenarios
    - viewing capacity usage trends 108
    - viewing information 225
    - viewing status 225
  - applications and departments scenarios
    - monitoring capacity usage 102

- architecture
  - description 31
- architecture, product 48
- audience ix
- Availability of Switch Ports report 635
- awareness, site 323, 412

## B

- backing up
  - considerations for 305
- balance pools 94
  - criteria 278
  - modify 278
  - volumes, colocate 101
  - volumes, re-tier 273, 275
- Balance Pools wizard 276
- Basic HyperSwap
  - about 356
  - session commands 384
- block storage provisioning
  - on vSphere Web Client extension 267
- block storage reports
  - viewing in vSphere Web Client extension 1025
  - vSphere Web Client extension 1025

## C

- capacity
  - adding to a capacity pool 252
  - removing from a capacity pool 254
  - volumes
    - Storage Virtualizer - Pools Space Report 523
    - Storage Virtualizer - Volumes Space Report 527
    - Storage Virtualizers - Space Report 517
- Capacity of One Group report 648
- capacity pool 70, 245
  - adding resources to 252
  - creating 252
  - deleting 254
  - modifying 254
  - removing resources from 254
  - viewing capacity of 253
  - viewing details of 253
  - viewing resources of 253
- CIM agents 88
- Cinder 1126, 1127
- Cinder driver 1126, 1127, 1135
- Cisco MDS 9000 87
- client ports 53
- cloud computing 1126, 1127
- cloud configuration 70, 245
- collection
  - overview 65

- columns
    - hiding in vSphere Web Client extension 1024
    - sorting in vSphere Web Client extension 1024
  - comments, sending xiii
  - communication-failure trap
    - descriptions 312
  - Compare Performance of Multiple Controllers and Modules report 694
  - Compare Performance of Multiple Disks report 703
  - Compare Performance of Multiple IO Groups report 715
  - Compare Performance of Multiple Managed Disks report 723
  - Compare Performance of Multiple Nodes report 695
  - Compare Performance of Multiple Pools report 735
  - Compare Performance of Multiple Ports report 742
  - Compare Performance of Multiple RAID Arrays report 748
  - Compare Performance of Multiple Storage Systems report 682
  - Compare Performance of Multiple Switch Ports report 646
  - Compare Performance of Multiple Switches report 639
  - Compare Performance of Multiple Volumes report 760
  - Compare Performance of One Pool over Time Ranges report 736
  - Compare Performance of One Storage System over Time Ranges report 683
  - Compare Performance of One Switch over Time Ranges report 640
  - Compare Performance of One Switch Port over Time Ranges report 647
  - components 31
  - computer alerts
    - triggering conditions 1062
  - concept diagram
    - VASA provider 79
  - configuration change alert
    - editing 163
  - configuration change alert for resource
    - editing 162
  - configuration change, SNMP trap
    - descriptions 311
  - configuring 70, 246
    - alerts for performance monitors 155
    - alerts for probes 149
    - logical paths 336
    - storage systems 329
  - connections
    - about 319
    - direct 320
    - Hardware Management Console (HMC) 321
    - modifying
      - host systems 334
    - modifying properties 328
    - z/OS 321
  - considerations, back up and recovery 305
  - consistency group 346
    - time interval 346
  - controls for performance view
    - export performance information 212
    - hide and show resources 212
    - hide and show the legend 212
    - remove resources 212
    - select metrics 212
    - specify a time range 212
    - specify granularity 212
    - view chart 212
    - view chart in a new window 212
    - view table 212
  - copy sets
    - about 341
    - adding 414
    - exporting 307, 431
    - importing 307, 431
    - removing 433
    - supported number of role pairs and volumes 341
    - viewing 476
  - creating 226, 227, 233, 246
    - FlashCopy session 414
    - Global Mirror session 422
    - Metro Global Mirror session 425
    - Metro Mirror session 419
    - Snapshot session 417
  - creating applications 226
  - creating departments 233
  - creating subcomponents 227
  - CSV file
    - exporting vSphere reports 1017
  - custom reports
    - See also* reports, custom
    - Hitachi Universal Storage Platform V systems 583
    - nodes 597
    - ports on Hitachi Universal Storage Platform V systems 620
    - storage virtualizer I/O groups 606, 607
    - storage virtualizer nodes 597
    - storage virtualizer pools 620, 625
    - storage virtualizer ports 615, 616
    - storage virtualizer volumes 587, 592
    - storage virtualizers 570, 574
    - storage-virtualizer managed disks 584, 586
    - volumes on Hitachi Universal Storage Platform V systems 595
  - customizing session lists 438
- D**
- dashboard 188
  - data collection 65
    - creating performance monitors in the web-based GUI 152
    - modifying performance monitors in the web-based GUI 154
    - performance monitors 151
    - starting and stopping performance monitors in the web-based GUI 156
    - verifying that a performance monitor is running in the web-based GUI 153
  - data collection (*continued*)
    - verifying that a probe is running in the web-based GUI 148
    - viewing all performance monitors in the web-based GUI 157
    - viewing performance monitor logs in the web-based GUI 158
    - viewing performance monitors for a resource in the web-based GUI 157
    - viewing performance monitors in the web-based GUI 157
  - data copying symbols 439
  - data exposure 346
  - data migration utility 434
  - data paths
    - changing the layout 243
    - exporting a data path to a file 244
    - fabrics 238
    - hypervisors 238
    - navigating a data path 242
    - overview 238
    - printing a view 244
    - resources that you can view 239
    - saving the view settings 244
    - servers 238
    - storage systems 238
    - topology view 239
    - topology view for data paths
      - changing the layout 243
      - exporting a data path to a file 244
      - navigating a data path 242
      - printing a view 244
      - saving the view settings 244
      - viewing information about internal resources 243
      - viewing information about top-level resources 242
      - zooming in 243
      - zooming out 243
    - viewing information about internal resources 243
    - viewing information about top-level resources 242
    - zooming in 243
    - zooming out 243
  - data store
    - creating in vSphere Web Client extension 267
  - datastore
    - See* data store
  - Datastores report
    - vSphere Web Client extension 1025, 1027
  - default
    - zone 86
  - department
    - scenarios
      - comparing storage in each department 104
  - department context 236
  - departments
    - applications 236
    - creating departments for monitoring 233
    - monitoring 233
    - removing 236

- departments (*continued*)
  - removing departments for monitoring 235
  - viewing information 233
  - viewing status 233
- departments and applications scenarios
  - different levels of business hierarchy 102
  - relational 108
- details 230, 235
- disabling
  - Metro Mirror heartbeat 430
  - Open HyperSwap 382
  - OpenSwap 382
- disaster recovery, practicing 434
- discover 159
- Disks Capacity report 696
- domain port zoning 86
- Domain\_ID 86

## E

- Easy Tier heat map transfer
  - about 318
  - configuring for 330
  - log file location 318
  - results, viewing 332
- EB
  - See* EiB
- EiB 83
- enabling, Metro Mirror heartbeat 430
- ESS/DS storage systems
  - adding logical paths 337
  - using a CSV file 337
  - removing logical paths 340
- ESX and ESXi
  - adding 136
- exported file systems 159
- exporting
  - copy sets 307, 431
  - historical data for Global Mirror role pairs 431
- exporting information in the web-based GUI 237, 285

## F

- fabric 85
- fabric alerts
  - triggering conditions 1062
- fabric connections
  - viewing in vSphere Web Client extension 1024
- Fabric Manager
  - zone 86
  - zone set 86
- fabrics
  - adding 133, 134
  - agent types for monitoring 134
  - modifying properties 281
  - viewing data paths 238
  - viewing information in the web-based GUI 185
  - viewing status in the web-based GUI 196

- failback 340, 347
- failover 340, 347
- FCID zoning 86
- file share
  - access path 251
- file storage provisioning
  - on vSphere Web Client extension 268
- file system reports
  - viewing in vSphere Web Client extension 1026
- File system reports
  - vSphere Web Client extension 1027
- File System to Volume Relationships report 633
- File Systems Capacity report 670
- file usage
  - volumes, tiering 97
- filtering lists in the web-based GUI 287
- filters 228
  - assign resources 228
  - sorting rows 288
- FlashCopy
  - about 357
  - creating a session 414
  - session commands 385

## G

- GiB 83
- Global Mirror
  - about 365
  - creating a session 422
  - session commands 391
  - session details, viewing 445
- Global Mirror with Practice, session commands 395
- glossary 1153
- granting
  - access privileges 480
- Groups Capacity report 650

## H

- hardware zoning 86
- health, viewing summary of 437
- heartbeat, enabling and disabling 430
- historical reports
  - Storage Virtualizer - Pools Historical Space Report 557
  - Storage Virtualizer - Volumes Historical Space Report 560
  - Storage Virtualizers - Historical Space Report 555
- host systems
  - adding a connection to 333
  - managing 333
  - modifying a connection to 334
  - removing 335, 336
  - viewing connection information for 477
- hypervisor alerts
  - triggering conditions 1054
- Hypervisor Data Stores Capacity report 659
- Hypervisor Disks Capacity report 660
- hypervisors 271

- hypervisors (*continued*)
  - adding 136
  - modifying properties 281
  - optimized volumes, colocate 101
  - tiering, analyze 271
  - viewing data paths 238
  - viewing information in the web-based GUI 185
  - viewing status in the web-based GUI 196
- Hypervisors Capacity report 652

## I

- I/O density
  - volumes, tiering 97
- I/O groups
  - performance statistics 542
  - Storage Virtualizers - Most Active IO Groups Report 530
- I/O rate
  - volumes, tiering 99
- IBM Tivoli Storage Productivity Center
  - architecture 31
  - component overview 31
  - overview 31
- IBM Tivoli Storage Productivity Center for Replication
  - starting 303
- icons
  - session status 439
  - storage systems status 317
- identifying the version of Tivoli Storage Productivity Center for Replication 305
- importing copy sets 307, 431
- Installation
  - Cinder driver package contents 1132
- interfaces
  - changes in V5.2, V5.2.1, V5.2.2 39
  - display differences 195
  - overview 33
  - when to use 34

## J

- journal volume, restoring data from 330

## K

- KB
  - See* KiB
- KiB 83

## L

- license restrictions 19
- licenses 18
- Linux 1126, 1127
- lists
  - hiding columns 289
  - reorganizing columns 289
  - showing columns 289
- locations
  - modifying for a site 432

- locations (*continued*)
  - modifying for a storage system 328
- logical paths
  - adding 337
    - using a CSV file 337
  - configuring 336
  - removing 340
  - viewing 337, 477
- logical storage subsystems
  - viewing volumes 477

## M

- managed disks
  - performance statistics 545
  - Storage Virtualizer - MDisk Details Report 544
  - Storage Virtualizers - Most Active MDisks Report 532
  - Storage Virtualizers - Most Active Nodes Report 533
  - Storage Virtualizers - Most Active Pools Report 535
  - Storage Virtualizers - Most Active Ports Report 538
  - Storage Virtualizers - Most Active Volumes Report 536
- Managed Disks Capacity report 521, 716
- management servers
  - about 308
  - reconnecting 315
  - setting as the standby server 313
  - setting up a remote server as the standby server 314
  - state change SNMP trap descriptions 312
  - taking over 315
  - viewing health summary 437
- managing
  - host systems 333
  - storage systems 317
- max deviation percentage
  - default range 278
- MB
  - See MiB
- MDisk Group Details report 787
- MDisk Group - VDisk Workload Activity Details report 783
- MDisk Groups - VDisk Workload Activity report 770
- MDisk Groups - Workload Activity report 790
- MDisks
  - performance statistics 545
  - Storage Virtualizer - MDisk Details Report 544
  - Storage Virtualizers - Most Active MDisks Report 532
- messages, viewing 477
- metricsperformance reports 1069
- Metro Global Mirror
  - about 370
  - creating a session 425
  - session commands 398
- Metro Global Mirror scenarios
  - disaster recovery (unplanned outages) 117

- Metro Global Mirror scenarios (*continued*)
  - planned outages 114
  - practice 120
- Metro Global Mirror with Practice, session commands 402
- Metro Mirror
  - about 358
  - creating a session 419
  - session commands 387, 406
- Metro Mirror heartbeat, enabling and disabling 430
- Metro Mirror with Practice, session commands 389
- MiB 83
- modifying 248
  - access privileges 481
  - host system connections 334
  - site location 432
  - storage connection properties 328
  - storage system location 328
- monitor
  - capacity and space usage 106
- monitoring 1100
- Monitoring Agent
  - displaying reports 89, 1028
- most active
  - I/O groups
    - storage virtualizers 530
  - managed disks
    - storage virtualizers 532
  - MDisks
    - storage virtualizers 532
  - nodes
    - storage virtualizers 533
  - pools
    - storage virtualizers 535
  - ports
    - storage virtualizers 538
  - Storage Virtualizers - Most Active MDisks Report 532
  - volumes
    - storage virtualizers 536
- Most Active Controllers or Modules report 685
- Most Active Disks report 698
- Most Active Host Connections report 704
- Most Active Hypervisors report 654
- Most Active IO Groups report 710
- Most Active Managed Disks report 718
- Most Active Nodes report 689
- Most Active Pools report 731
- Most Active Ports report 738
- Most Active RAID Arrays report 743
- Most Active Servers report 664
- Most Active Storage Systems report 678
- Most Active Switch Ports report 641
- Most Active Volumes report 756
- multiple resources alert definition
  - viewing 172

## N

- N\_Port address identifier 86
- N\_Port\_Name 86
- NAS mounts report
  - vSphere Web Client extension 1027

- Node\_Name 86
- nodes
  - most active 534
  - performance statistics 547
  - Storage Virtualizer - Node Details Report 546
  - Storage Virtualizers - Most Active Nodes Report 533
- non standard zones
  - fabric assist 88
  - protocol 88
  - quick loop 88
- notifications for alerts 1065

## O

- Open HyperSwap
  - disabling 382
- open source 1126, 1127
- Opening the management GUI for a resource 284
- OpenStack 1126, 1127
  - Installing Tivoli Storage Productivity Center 1133
  - Tivoli Storage Productivity Center 1129
  - Uninstalling Tivoli Storage Productivity Center Cinder driver 1134
- OpenSwap
  - disabling 382
- overview 70

## P

- paths
  - adding 337
    - using a CSV file 337
  - removing 340
- PB
  - See PiB
- performance 1069
  - controls for the performance view 212
  - hot spots 93
  - metrics
    - disk 1069
    - port 1069
    - volume 1069
  - monitoring 205
  - performance view 211, 218
  - storage virtualizer reports
    - managed disks 544
    - pools 553
  - Storage Virtualizers - Most Active IO Groups Report 531
  - viewing alerts for threshold violations 220
  - viewing thresholds 218, 219
- performance alert for resource
  - editing 164, 166
  - suppressing notifications 169
- performance chart
  - viewing in vSphere Web Client extension 1017
  - vSphere Web Client extension 1019



- performance data
  - analysis period, change 281
- Performance Data Export report 684
- Performance of One Controller or Module report 687
- Performance of One Disk report 700
- Performance of One IO Group report 712
- Performance of One Managed Disk report 720
- Performance of One Node report 691
- Performance of One Pool report 733
- Performance of One Port report 740
- Performance of One RAID Array report 745
- Performance of One Storage System report 680
- Performance of One Switch Port report 643
- Performance of One Switch report 637
- Performance of One Volume report 758
- Performance of Volumes by Host Connection report 709
- Performance of Volumes by Hypervisor report 657
- Performance of Volumes by Server report 667
- performance reporting
  - Hitachi Universal Storage Platform V systems 583
  - nodes 597
  - ports on Hitachi Universal Storage Platform V systems 620
  - storage virtualizer I/O groups 606, 607
  - storage virtualizer nodes 597
  - storage virtualizer pools 620, 625
  - storage virtualizer ports 615, 616
  - storage virtualizer volumes 587, 592
  - storage virtualizers 570, 574
  - storage-virtualizer managed disks 584, 586
  - volumes on Hitachi Universal Storage Platform V systems 595
- performance statistics
  - volumes 551
- performance thresholds 1100
- performance view
  - accessing 211, 218
  - bookmarks 218
  - controls 212
  - favorites 218
  - saving URLs 218
  - viewing the performance of internal resources 216
  - viewing the performance of related resources 216
- performing
  - takeover 315
- PiB 83
- planned outage of remote (H3) and return 117
- planned outage of the local site (H1) while in normal configuration, with a production move to the remote site (H3) 116
- planned outage of the local site (H1), with a production move to the intermediate site (H2) and an optional return 114
- planned outage on site 2 (H2) and return three sites to the original configuration 115
- pools
  - balance 94, 95, 273
  - balance, analyze 276
  - hot spots, identify 94
  - modifying properties 281
  - most active 536
  - performance statistics 549
  - space usage 558
  - Storage Virtualizer - Pool Details Report 548
  - Storage Virtualizer - Pools Space Report 523
  - Storage Virtualizers - Most Active Pools Report 535
  - tier 95
  - tiering, analyze 272
  - volumes, re-tier 273
- Pools Capacity report 724
- Pools Historical Capacity report 565, 728
- port name zoning 86
- port zoning 86
- PortId zoning 86
- ports 53
  - most active 539
  - performance statistics 554
  - Storage Virtualizers - Most Active Ports Report 538
- practice at H3 for Metro Global Mirror 123
- practice recover to H1 with production continuing on H2 121
- practice recover to H2 with production continuing on H1 in a normal configuration 120
- practice volumes 346, 435
- practicing disaster recovery 434
- predefined reports
  - historical information 555
  - Storage Virtualizer - IO Group Details Report 542
  - Storage Virtualizer - MDisk Details Report 544, 545
  - Storage Virtualizer - Node Details Report 546, 547
  - Storage Virtualizer - Performance Details Report 540
  - Storage Virtualizer - Pool Details Report 548, 549
  - Storage Virtualizer - Pools Historical Space Report 557, 558
  - Storage Virtualizer - Pools Space Report 523
  - Storage Virtualizer - Port Details Report 553, 554
  - Storage Virtualizer - Volume Details Report 551
  - Storage Virtualizer - Volumes Historical Space Report 560, 561
  - Storage Virtualizer - Volumes Space Report 527
- predefined reports (*continued*)
  - Storage Virtualizer Performance Details Report 540
  - Storage Virtualizers - Historical Space Report 555, 556
  - Storage Virtualizers - Most Active IO Groups Report 530, 531
  - Storage Virtualizers - Most Active MDisks Report 532
  - Storage Virtualizers - Most Active Nodes Report 533, 534
  - Storage Virtualizers - Most Active Pools Report 535, 536
  - Storage Virtualizers - Most Active Ports Report 538, 539
  - Storage Virtualizers - Most Active Volumes Report 536, 537
  - Storage Virtualizers - Space Report 517
- Preserve Mirror option, recommendations for 413
- probes
  - configuring alerts 149
- problems
  - report xiv
- product architecture 48
- properties
  - session
    - SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified sessions 471
    - TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions 447
    - XIV system sessions 474
    - viewing session 447
- protected volumes, about 323
- protecting volumes 329
- Provision Storage wizard 262
- provisioning 159, 262
  - custom requirements for 249
  - overview 75, 245
  - scenarios
    - constrained set of storage 112
    - using vSphere Web Client extension 112
  - shares 265, 266
  - volumes 263, 264

## R

- re-tiering volumes 275
- read I/O capability
  - calculating for MDisk groups 795
  - calculating for pools 795
  - formula 795
- reader feedback, sending xiii
- recommendations for using the Preserve Mirror option 413
- reconnecting the active and standby management servers 315
- recovering, considerations for 305
- Red Hat 1126, 1127
- refreshing, storage system configuration 329

- removing 231, 235
  - copy sets 433
  - host systems 335, 336
  - logical paths 340
  - sessions 433
  - storage systems 145, 327
  - user access 482
- removing applications 231
- removing departments 235
- removing resources 144
- report displays
  - customize in vSphere Web Client extension 1024
- reporting
  - trcmd** command
    - running reports 502
    - saving report output to file system 502
  - Cognos
    - administrative privileges, restricting 511
    - creating and viewing reports within 494
    - password change 1009
    - timeout 1007
  - Cognos software
    - configuring security 509
  - archive location, specifying 503
  - backup copies of reports, creating 506
  - build ID, showing 1013
  - CM-REQ-4069 error message 1011
  - CM-REQ-4158 error message 1011
  - Cognos
    - See* Tivoli Common Reporting
  - configuring, overview 498
  - connections to TPCDB, multiple 505
  - copying a predefined report 500
  - creating a folder 506
  - creating custom reports 493
  - creating reports from within Cognos 494
  - custom reports
    - See* reports, custom
  - dashboard 188
  - DPR-ERR-2109 error message 1010
  - email
    - configuring to send reports by 501
    - sending report by 501
  - file system
    - configuring report to save to 504
    - saving reports to 503
    - specifying archive location 503
    - specifying locations for output 504
    - specifying root directory 503
  - filters, limit of 5000 items 1011
  - format and delivery options, setting 498
  - getting started 483
  - IBM Cognos Configuration GUI, starting 500
  - logo, customizing 506
  - long time to complete performance report 1012
  - long time to display report 1011
- reporting (*continued*)
  - No data available message 1010
  - package version number, showing 1013
  - PDF output, maximum number of columns 1008
  - Portuguese (Brazilian) text in reports 1012
  - predefined reports
    - See* reports, predefined
  - Query Studio, performance report 1012
  - remote connection, logging off 1011
  - report contains no data 1011
  - report titles
    - customizing 506
    - modifying 507
  - report views 499
  - scenarios
    - analyzing hypervisor volumes 495
    - identifying relationships 498
    - investigating pool performance 497
    - monitoring volume performance 496
  - scheduling a report with predefined options 499
  - search index, creating 507
  - security
    - access to functions, restricting 514
    - access to reports, specifying 513
    - administrative privileges, assigning 510
    - administrative privileges, restricting 511
    - configuring for Cognos software 509
    - roles 510, 511
    - user IDs, creating 509
    - WebSphere Application Server, user IDs 509
  - Storage Volume properties 1010
  - Tivoli Common Reporting
    - See also* reporting, trcmd command
    - access to functions, restricting 514
    - access to reports, specifying 513
    - administrative privileges, assigning 510
    - opening 510
    - relationship to Cognos 483
    - timeout 1007
  - TPCDB
    - multiple connections to 505
    - password change 1009
    - password prompt 1009
  - troubleshooting 1007
  - viewing predefined reports 493
  - viewing reports from within Cognos 494
- reporting problems xiv
- reports
  - storage tier 761
  - Storage Virtualizer - IO Group Details Report 542
  - Storage Virtualizer - MDisk Details Report 544
- reports (*continued*)
  - Storage Virtualizer - Node Details Report 546, 547
  - Storage Virtualizer - Performance Details Report 540
  - Storage Virtualizer - Pool Details Report 548, 549
  - Storage Virtualizer - Pools Historical Space Report 558
  - Storage Virtualizer - Pools Space Report 523
  - Storage Virtualizer - Port Details Report 553, 554
  - Storage Virtualizer - Volume Details Report 551
  - Storage Virtualizer - Volumes Historical Space Report 561
  - Storage Virtualizer - Volumes Space Report 527
  - Storage Virtualizer - Volume Details Report 551
  - Storage Virtualizer Performance Details Report 540
  - Storage Virtualizers - Historical Space Report 556
  - Storage Virtualizers - Most Active IO Groups Report 530, 531
  - Storage Virtualizers - Most Active MDisks Report 532
  - Storage Virtualizers - Most Active Nodes Report 533, 534
  - Storage Virtualizers - Most Active Pools Report 535, 536
  - Storage Virtualizers - Most Active Ports Report 538, 539
  - Storage Virtualizers - Most Active Volumes Report 536, 537
  - Storage Virtualizers - Space Report 517
  - threshold values 798
  - upgrade from Version 5.1, Version 5.1.1, or Version 5.1.1.1 to Version 5.2 (or later) 516
  - vSphere 1013
  - vSphere Web Client extension 1013
  - reports, custom
    - block server storage systems, performance 823
    - capacity and relationships reports
      - creating 927
      - overview 927
    - clusters
      - hypervisors, capacity 945
      - servers, capacity 933
    - controllers, storage systems
      - block server systems, performance 862
      - DS8000 and DS6000, performance 858
    - creating 493
    - data stores on hypervisors, capacity 946
    - DS8000 and DS6000 storage systems, performance 813
    - fabrics, capacity 1006
    - file system groups on servers, capacity 942

- reports, custom *(continued)*
  - file systems
    - hypervisors, capacity 945
    - NAS systems, capacity 952
    - servers, capacity 936
  - filters
    - adding to capacity and relationships reports 929
    - hypervisors 930, 932
    - limit of 5000 items 1011
    - servers 929
    - storage resource groups 933
    - storage systems 932
  - groups, capacity 1002
  - help 1009
  - host connections, back-end storage
    - SAN Volume Controller, performance 908
    - Storwize V7000 Unified, performance 908
    - Storwize V7000, performance 908
  - host connections, storage systems
    - DS8000 and DS6000, performance 911
    - XIV systems, performance 915
  - hypervisor controllers, capacity 949
  - hypervisor multipath drivers, capacity 949
  - hypervisors, capacity 943
  - I/O groups
    - SAN Volume Controller, performance 896
    - Storwize V7000 Unified, performance 896
    - Storwize V7000, performance 896
  - IBM SONAS 958
    - disks, capacity 955
    - pools, capacity 957
  - local disks, storage systems
    - capacity 978
    - performance 905, 906
  - logical volumes
    - NAS systems, capacity 955
    - servers, capacity 937
  - managed disks
    - SAN Volume Controller, performance 884
    - Storwize V7000 Unified, performance 884
    - Storwize V7000, performance 884
  - managed disks, storage systems
    - capacity 1000
    - performance 882
  - managed disks, storage virtualizers
    - capacity 976
  - modules, storage systems
    - XIV systems, performance 864
  - NAS disks, capacity 955
  - NAS exports, capacity 954
  - NAS systems, capacity 951
  - NetApp disks, capacity 955
  - No data available message 1010
  - nodes
    - SAN Volume Controller, performance 887
    - Storwize V7000 Unified, performance 887

- reports, custom *(continued)*
  - nodes *(continued)*
    - Storwize V7000, performance 887
  - overview 491
  - performance reports
    - creating 799
    - overview 798
  - pools
    - SAN Volume Controller, performance 873
    - Storwize V7000 Unified, performance 873
    - Storwize V7000, performance 873
  - pools, storage systems
    - DS8000 and DS6000, performance 875
    - XIV systems, performance 880
  - ports
    - SAN Volume Controller, performance 843
    - Storwize V7000 Unified, performance 843
    - Storwize V7000, performance 843
  - ports, storage systems
    - block server systems, performance 851
    - DS8000 and DS6000, performance 846
    - XIV systems, performance 850
  - SAN Volume Controller, performance 804
  - server controllers, capacity 941
  - server disks on hypervisors, capacity 947
  - server disks, capacity 939
  - server groups, capacity 940
  - server multipath drivers, capacity 942
  - servers, capacity 933
  - storage host connections, performance 908
  - storage resource groups, capacity 1002
  - storage system controllers, performance 858
  - storage system disks, capacity 998
  - storage system groups, capacity 1000
  - storage system I/O groups, performance 895
  - storage system modules, performance 863
  - storage system nodes, performance 886
  - storage system pools, capacity 984
  - storage system pools, performance 867
  - storage system ports, performance 841
  - storage system RAID arrays, performance 852, 853
  - storage system volumes, capacity 990
  - storage system volumes, performance 824
  - storage systems, capacity 980
  - storage systems, performance 800

- reports, custom *(continued)*
  - storage virtualizer groups, capacity 980
  - storage virtualizer pools, capacity 963
  - storage virtualizer volumes, capacity 969
  - storage virtualizers, capacity 959
  - Storage Volume properties 1010
  - Storwize V7000 Unified
    - disks, capacity 955
    - filesets, capacity 958
    - pools, capacity 957
  - Storwize V7000 Unified, performance 804
  - Storwize V7000, performance 804
  - switch ports, capacity 1005
  - switch ports, performance 922, 923
  - switches, capacity 1003
  - switches, performance 918, 919
  - virtual machines managed by hypervisors, capacity 949
  - volume groups on servers, capacity 938
  - volumes
    - SAN Volume Controller, performance 830
    - Storwize V7000 Unified, performance 830
    - Storwize V7000, performance 830
  - volumes, storage systems
    - block server systems, performance 840
    - DS8000 and DS6000, performance 834
    - XIV systems, performance 837
  - XIV systems, performance 820
- reports, predefined
  - Availability of Switch Ports 635
  - Capacity of One Group 648
  - Compare Performance of Multiple Controllers and Modules 694
  - Compare Performance of Multiple Disks 703
  - Compare Performance of Multiple IO Groups 715
  - Compare Performance of Multiple Managed Disks 723
  - Compare Performance of Multiple Nodes 695
  - Compare Performance of Multiple Pools 735
  - Compare Performance of Multiple Ports 742
  - Compare Performance of Multiple RAID Arrays 748
  - Compare Performance of Multiple Storage Systems 682
  - Compare Performance of Multiple Switch Ports 646
  - Compare Performance of Multiple Switches 639
  - Compare Performance of Multiple Volumes 760
  - Compare Performance of Multiple Volumes, scenario 496

- reports, predefined (*continued*)
  - Compare Performance of One Pool over Time Ranges 736
  - Compare Performance of One Pool over Time Ranges, scenario 497
  - Compare Performance of One Storage System over Time Ranges 683
  - Compare Performance of One Switch over Time Ranges 640
  - Compare Performance of One Switch Port over Time Ranges 647
  - copying 500
  - Disks Capacity 696
  - File System to Volume Relationships 633
  - File Systems Capacity 670
  - Groups Capacity 650
  - Hypervisor Data Stores Capacity 659
  - Hypervisor Disks Capacity 660
  - Hypervisors Capacity 652
  - Managed Disks Capacity 521, 716
  - MDisk Group Details 787
  - MDisk Group - VDisk Workload Activity Details 783
  - MDisk Groups - VDisk Workload Activity 770
  - MDisk Groups - Workload Activity 790
  - Most Active Controllers or Modules 685
  - Most Active Disks 698
  - Most Active Host Connections 704
  - Most Active Hypervisors 654
  - Most Active Hypervisors, scenario 495
  - Most Active IO Groups 710
  - Most Active Managed Disks 718
  - Most Active Nodes 689
  - Most Active Pools 731
  - Most Active Ports 738
  - Most Active RAID Arrays 743
  - Most Active Servers 664
  - Most Active Storage Systems 678
  - Most Active Switch Ports 641
  - Most Active Volumes 756
  - overview 484
  - Performance Data Export report 684
  - performance metrics, searching for 494
  - Performance of One Controller or Module 687
  - Performance of One Disk 700
  - Performance of One IO Group 712
  - Performance of One Managed Disk 720
  - Performance of One Node 691
  - Performance of One Pool 733
  - Performance of One Port 740
  - Performance of One RAID Array 745
  - Performance of One Storage System 680
  - Performance of One Switch 637
  - Performance of One Switch Port 643
  - Performance of One Volume 758
  - Performance of Volumes by Host Connection 709

- reports, predefined (*continued*)
  - Performance of Volumes by Hypervisor 657
  - Performance of Volumes by Hypervisor, scenario 495
  - Performance of Volumes by Server 667
  - Pools Capacity 724
  - Pools Historical Capacity 565, 728
  - properties, searching for 494
  - report views 499
  - scenarios
    - analyzing hypervisor volumes 495
    - identifying relationships 498
    - investigating pool performance 497
    - monitoring volume performance 496
  - scheduling with predefined options 499
  - security, specifying access to reports 513
  - Server Disks Capacity 669
  - Servers Capacity 662
  - Storage Resource Group - VDisk Workload Activity 779
  - Storage Resource Group - VDisk Workload Activity Details 763
  - Storage Resource Relationships Summary 628
  - Storage Resource Relationships Summary (Configurable) 630
  - Storage Resource Relationships Summary, scenario 498
  - Storage Systems Capacity 672
  - Storage Systems Historical Capacity 563, 675
  - Storage Virtualizer - VDisk Workload Activity 774
  - Summarized Performance of Volumes by Host Connection 706
  - Summarized Performance of Volumes by Hypervisor 655
  - Summarized Performance of Volumes by Hypervisor, scenario 495
  - Summarized Performance of Volumes by Server 665
  - VDisk Details 767
  - viewing descriptions of 493
  - viewing predefined reports 493
  - Volumes Capacity 749
  - Volumes Historical Capacity 568, 754
- resource alert definition viewing 170
- resources
  - adding 127
  - adding agentless servers for monitoring 139
  - adding fabrics for monitoring 133
  - adding hypervisors for monitoring 136, 137
  - adding storage systems for monitoring 132
  - adding switches for monitoring 133
  - adding to a capacity pool 252

- resources (*continued*)
  - manually adding resources to applications 229
  - modifying properties 281
  - removing 144
  - removing from a capacity pool 254
  - Storage Virtualizer - Pools Space Report
    - capacity 523
    - run 523
    - used and available space 523
  - Storage Virtualizer - Volumes Space Report
    - capacity 527
    - run 527
    - used and available space 527
  - Storage Virtualizers - Space Report
    - capacity 517
    - run 517
    - used and available space 517
  - restoring data from a journal volume 330
  - restrictions
    - based on license 19
    - based on role 19
  - role pairs 346
  - role pairs, status messages 444
  - role restrictions 19
  - roles 83
  - run
    - Storage Virtualizer - Pools Historical Space Report 557
    - Storage Virtualizer - Volumes Historical Space Report 560
    - Storage Virtualizers - Historical Space Report 555

## S

- SAN Volume Controller publications ix
- scenarios
  - overview of scenarios for practicing at H3 for Metro Global Mirror 123
    - practice scenario 1 124
    - practice scenario 2 124
    - practice scenario 3 124
  - planned outage of remote (H3) and return 117
  - planned outage of the local site (H1) with a production move to the remote site (H3) 116
  - planned outage of the local site (H1), with a production move to the intermediate site (H2) and an optional return 114
  - planned outage on site 2 (H2) and return three sites to the original configuration 115
  - practice at H3 for Metro Global Mirror 123
  - practice recover to H1 with production continuing on H2 121
  - practice recover to H2 with production continuing on H1 in a normal configuration 120

- scenarios (*continued*)
    - unplanned outage of local (H1) with production move to remote (H3) 119
    - Unplanned outage of local (H1), production move to intermediate (H2) 117
    - unplanned outage of remote (H3) 119
    - unplanned outage of Site 2 (H2) 118
  - scenarios for Metro Global Mirror 114, 117, 120
  - scenarios overview for practicing at H3 for Metro Global Mirror 123
    - practice scenario 1 124
    - practice scenario 2 124
    - practice scenario 3 124
  - SCSI Volumes (LUNs) report
  - vSphere Web Client extension 1025
  - security 478
    - Cognos software, configuring 509
  - sending comments xiii
  - server alerts
    - triggering conditions 1062
  - Server Disks Capacity report 669
  - servers
    - adding
      - importing servers from a list 141
      - importing servers with agents from a file 143
    - adding by deploying an agent 142
    - adding without agents 139
    - modifying properties 281
    - optimized volumes, colocate 101
    - tiering, analyze 271
    - viewing data paths 238
    - viewing information in the web-based GUI 185
    - viewing status in the web-based GUI 196
  - Servers Capacity report 662
  - service xiii
  - service class 70, 245, 246, 248
    - deleting 249
    - satisfying custom requirements of a 249
    - viewing properties of 247
    - viewing storage constraints of 247
    - viewing users who can provision by using 247
    - viewing volumes or shares associated with 247
  - service management connect xiii
  - session commands
    - Basic HyperSwap 384
    - FlashCopy 385
    - Global Mirror 391
    - Global Mirror with Practice 395
    - Metro Global Mirror 398
    - Metro Global Mirror with Practice 402
    - Metro Mirror 387, 406
    - Metro Mirror with Practice 389
    - Snapshot 386
  - session state change trap
    - descriptions 310
  - session types
    - Basic HyperSwap 356
    - FlashCopy 357
    - Global Mirror 365
    - Metro Global Mirror 370
    - Metro Mirror 358
    - Snapshot 357
  - sessions
    - about 340
    - customizing lists of 438
    - FlashCopy 414
    - Global Mirror 422
    - images 439
    - Metro Global Mirror 425
    - Metro Mirror 419
  - properties
    - SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified sessions 471
    - TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions 447
    - XIV system sessions 474
  - removing 433
  - Snapshot 417
  - states 442
  - status icons 439
  - viewing 438
  - viewing details 427, 445
    - Global Mirror 445
  - viewing health summary 437
  - viewing properties 447
- setting
  - remote standby management server 314
  - SNMP 316
  - standby management server 313
  - volume protection 329
- shares 159
- single sign-on
  - overview 90
  - user authentication method used with 90
- site awareness 323, 412
- site, modifying location of 432
- SMC xiii
- Snapshot
  - about 357
  - creating a session 417
  - session commands 386
- SNMP
  - adding managers 316
  - setting up 316
  - viewing alerts 438
- SNMP trap descriptions
  - communication-failure 312
  - configuration change 311
  - management servers state change 312
  - session state change 310
  - suspending-event notification 312
- software zoning 86
- sorting rows in lists 288
- space usage
  - storage virtualizers 556
- standby management servers
  - about 308
  - ports 53
  - setting 313
  - setting a remote server 314
- starting
  - IBM Tivoli Storage Productivity Center for Replication 303
  - Tivoli Storage Productivity Center 42
  - Tivoli Storage Productivity Center for Replication 45
  - Tivoli Storage Productivity Center stand-alone GUI 42
  - Tivoli Storage Productivity Center web-based GUI 44
- states, session 442
- status messages, role pairs 444
- status of resources in the web-based GUI
  - acknowledging the status of a resource 202
  - acknowledging the status of a top-level resource 202
  - acknowledging the status of an internal resource 204
  - how the status of resources is determined 197
  - viewing the overall status of resources 200
  - viewing the status of internal resources 201
  - viewing the status of top-level resources 201
- status, session icons 439
- stopping
  - Tivoli Storage Productivity Center for Replication 304
- storage connections
  - viewing details 476
- storage mapping
  - viewing in vSphere Web Client extension 1014
  - vSphere Web Client extension 1015
- Storage mapping
  - vSphere Web Client extension 1017
- storage pools
  - tier level, setting 270
  - tiering, analyze 272
  - volumes, re-tier 275
- storage provisioning
  - vSphere Web Client extension 80
- storage reporting
  - vSphere Web Client extension 80
- Storage Resource Group - VDisk
  - Workload Activity Details report 763
  - Workload Activity report 779
- storage resource groups
  - volumes space usage 561
- Storage Resource Relationships Summary (Configurable) report 630
- Storage Resource Relationships Summary report 628
- storage system alerts
  - triggering conditions 1037
- storage system metrics
  - customizing report displays in vSphere Web Client extension 1024

- storage system metrics (*continued*)
  - hiding columns 1024
  - sorting columns 1024
  - viewing in vSphere Web Client extension 1017
  - vSphere Web Client extension 1019
- Storage system metrics vSphere Web Client extension vSphere Web Client extension
  - Exporting information to a CSV file 1017
- storage systems
  - about 317
  - adding 132
  - adding a connection to
    - DS6000 326, 327
    - DS8000 324, 325
    - SAN Volume Controller 326
    - Storwize 326
    - TotalStorage Enterprise Server 326, 327
    - z/OS 327
  - icons 317
  - managing 317
  - modifying connection properties 328
  - modifying location of 328
  - modifying properties 281
  - ports 53
  - refreshing the configuration 329
  - removing 145, 327
  - supported 317
  - viewing data paths 238
  - viewing details 476
  - viewing health summary 437
  - viewing information in the web-based GUI 185
  - viewing status in the web-based GUI 196
  - viewing volumes 477
- Storage Systems Capacity report 672
- Storage Systems Historical Capacity report 563, 675
- storage tier reports
  - threshold values 798
  - threshold values, change 797
- storage tiering
  - pools, balancing 273
- storage tiers
  - workload activity reports 761
- Storage Virtualizer - IO Group Details Report
  - performance statistics 542
- Storage Virtualizer - MDisk Details Report
  - performance statistics 545
  - run 544
- Storage Virtualizer - Node Details Report
  - performance statistics 547
  - run 546
- Storage Virtualizer - Performance Details Report
  - performance statistics 540
- Storage Virtualizer - Pool Details Report
  - performance metrics 549
  - run 548
- Storage Virtualizer - Pools Historical Space Report
  - space usage 558
- Storage Virtualizer - Pools Space Report
  - capacity 523
  - capacity and space report 523
  - run 523
  - used and available space 523
- Storage Virtualizer - Port Details Report
  - performance statistics 554
  - run 553
- Storage Virtualizer - Volume Details Report
  - run 551
- Storage Virtualizer - Volumes Historical Space Report
  - space usage 561
- Storage Virtualizer - Volumes Space Report
  - capacity 527
  - capacity and space report 527
  - run 527
  - used and available space 527
- Storage Virtualizer - IO Group Details Report
  - run 542
- Storage Virtualizer - Volume Details Report
  - performance statistics 551
  - predefined reports 551
- Storage Virtualizer Performance Details Report
  - run 540
- Storage Virtualizer - VDisk Workload Activity report 774
- storage virtualizers 271, 570
  - custom reports
    - I/O groups 606, 607
    - managed disks 584, 586
    - nodes 597
    - pools 620, 625
    - ports 615, 616
    - storage virtualizers 570, 574
    - volumes 587, 592
  - disks 517
  - most active reports 530
  - performance details 540
  - pools 517
  - pools, balance 95
  - space usage 556
  - Storage Virtualizer - IO Group Details Report 542
  - Storage Virtualizer - Node Details Report
    - nodes 546
  - Storage Virtualizer - Performance Details Report 540
  - Storage Virtualizer - Pool Details Report 548
  - Storage Virtualizer Performance Details Report 540
  - Storage Virtualizers - Most Active IO Groups Report 530
  - Storage Virtualizers - Most Active MDisks Report 532
  - Storage Virtualizers - Most Active Nodes Report 533
- storage virtualizers (*continued*)
  - Storage Virtualizers - Most Active Pools Report 535
  - Storage Virtualizers - Most Active Ports Report 538
  - Storage Virtualizers - Most Active Volumes Report 536
  - Storage Virtualizers - Space Report 517
  - tiering, analyze 272
  - virtual volumes 517
  - volumes, analyze placement 93
  - volumes, optimize 93
  - volumes, tier 95
- Storage Virtualizers - Historical Space Report
  - space usage 556
- Storage Virtualizers - Most Active IO Groups Report
  - most active 531
  - running 530
- Storage Virtualizers - Most Active MDisks Report
  - most active 532
  - running 532
- Storage Virtualizers - Most Active Nodes Report
  - performance statistics 534
  - running 533
- Storage Virtualizers - Most Active Pools Report
  - performance statistics 536
  - running 535
- Storage Virtualizers - Most Active Ports Report
  - performance statistics 539
  - running 538
- Storage Virtualizers - Most Active Volumes Report 537
  - running 536
- Storage Virtualizers - Space Report
  - capacity 517
  - capacity and space 517
  - run 517
  - used and available space 517
- storage virtualizers reports
  - view 517
- subcomponents
  - creating subcomponents for existing applications 227
- subdepartments
  - adding subdepartments for monitoring 234
- Summarized Performance of Volumes by Host Connection report 706
- Summarized Performance of Volumes by Hypervisor report 655
- Summarized Performance of Volumes by Server report 665
- support xiii
- Support
  - contact xiv
- suspending-event notification trap
  - descriptions 312
- switch alerts
  - triggering conditions 1056

- switches
  - adding 133, 134
  - agent types for monitoring 134
  - modifying properties 281
  - viewing information in the web-based GUI 185
  - viewing status in the web-based GUI 196
- symbols
  - data copying 439
  - session status 439
  - volume roles 439

## T

- taking over management from the active server 315
- tasks
  - analysis execution 298, 299
  - creating provisioning tasks in the web-based GUI 294
  - creating tiering storage, balancing pools, and transforming-storage tasks in the web-based GUI 296
  - deleting in the web-based GUI 301
  - exporting provisioning tasks in the web-based GUI 295
  - implementing optimization
    - recommendations in the web-based GUI 298
  - managing in the web-based GUI 289
  - managing provisioning tasks in the web-based GUI 293
  - managing tiering storage, balancing pools, and transforming-storage tasks in the web-based GUI 295
  - pausing, resuming, and canceling analysis-execution tasks in the web-based GUI 300
  - running provisioning tasks in the web-based GUI 294
  - running tiering analysis, transform-plan, and balance-analysis tasks in the web-based GUI 297
  - scheduling the implementation of optimization recommendations in the web-based GUI 299
  - setting the number of task runs displayed 293
  - viewing all tasks in the web-based GUI 290
  - viewing details in the web-based GUI 291
  - viewing in the web-based GUI 290
  - viewing pending tasks 290
  - viewing task logs in the web-based GUI 292
  - viewing tasks for a resource in the web-based GUI 291
  - viewing tasks for a resource type in the web-based GUI 290

## TB

See TiB

- temporary zone sets 86
- threshold
  - alert 110
  - array 1102

- threshold (*continued*)
  - array thresholds 1102
  - controller 1110
  - controller thresholds 1110
- threshold values
  - storage tier reports, change 797
- thresholds 1100
  - port 1112
  - port thresholds 1112
  - viewing 218, 219
  - viewing alerts 220
- TiB 83
- tier levels
  - storage pools, setting 270
- tiering
  - analysis period 281
  - file usage 273
  - I/O density 273
  - I/O rate 273
- tiering storage
  - file usage 275
  - I/O density 275
  - I/O rate 275
- tiering thresholds
  - files last accessed, modify 274
  - I/O rate, modify 274
- Tivoli Storage Productivity Center
  - Cinder driver requirements 1131
  - VASA provider 79
- Tivoli Storage Productivity Center for Replication
  - stopping 304
- Tivoli Storage Productivity Center
  - limitations and known issues 1136
  - topology view for data paths 239
- TPCD.cfg
  - saveNonRoot parameter, add 159
- trademarks 1151
- translations
  - browser locale requirement xii
- triggering actions 1065
- triggering conditions
  - alerts 1036
  - computer alerts 1062
  - fabric alerts 1062
  - hypervisor alerts 1054
  - server alerts 1062
  - storage system alerts 1037
  - switch alerts 1056
- troubleshooting
  - performance 205
  - reports 1007

## U

- Ubuntu 1126, 1127
- units of measurement 83
- unplanned outage of local (H1) with production move to remote (H3) 119
- Unplanned outage of local (H1), production move to intermediate (H2) 117
- unplanned outage of remote (H3) 119
- unplanned outage of Site 2 (H2) 118
- user administration 478
  - granting access privileges 480
  - modifying access privileges 481

- user administration (*continued*)
  - removing user access 482
  - viewing access privileges 481
- user interfaces
  - changes in V5.2, V5.2.1, V5.2.2 39
  - when to use 34
- users
  - about 478
  - assigning roles 83
  - groups 478
  - roles 479
- using 228
- using practice volumes 434

## V

- VASA provider
  - architecture 79
  - Tivoli Storage Productivity Center 79
- VASA provider for Tivoli Storage Productivity Center
  - architecture 79
- vCenter Server system
  - adding 136, 137
- VDisk Details report 767
- ver command 305
- version of Tivoli Storage Productivity Center for Replication 305
- viewing
  - access privileges 481
  - copy sets 476
  - error messages 477
  - health summary 437
  - host system connection details 477
  - logical paths 337, 477
  - scenarios
    - connectivity information 113
    - NPIV connections 113
  - session details 427, 445
    - Global Mirror 445
  - session properties 447
  - sessions 438
  - SNMP alerts 438
  - storage connection details 476
  - storage system details 476
  - Tivoli Storage Productivity Center
    - VASA provider reports 1027, 1028
  - VASA provider reports 1027
  - viewing details of a department 235
  - viewing details of an application 230
  - volume details 477
- Viewing
  - the overall status of a storage environment 188
- viewing application details 230
- viewing department details 235
- viewing in the vSphere Web Client
  - VASA provider reports 1028
- viewing in vSphere Client
  - VASA provider reports 1027
- VMware vSphere
  - provisioning 112
- volume collocation 101
- volume protection
  - about 323
  - setting 329

- volume roles
  - about 344
  - symbols 439
- volumes
  - most active 537
  - practice 346, 435
  - role pairs 346
  - roles 344
  - Storage Virtualizer - Volume Details Report 551
  - Storage Virtualizer - Volumes Space Report 527
  - Storage Virtualizers - Most Active Volumes Report 536
  - tiering 97, 99
  - tiering by file usage 273
  - tiering by I/O density 273
  - tiering by I/O rate 273
  - tiering, analyze 272
  - viewing details 477
  - workload 276
- Volumes Capacity report 749
- Volumes Historical Capacity report 568, 754
- volumesstorage resource groups
  - storage virtualizers 561
- VSANs 87
- vSphere
  - reports 1013, 1014, 1015, 1017, 1019, 1024, 1025, 1026, 1027
- vSphere reports
  - exporting to a CSV file 1017
- vSphere resources
  - managing 266, 267, 268
- vSphere storage
  - provisioning 80
- vSphere storage reports
  - viewing 80
- vSphere Web Client extension
  - block storage reports 1025
  - datastores report 1027
  - Datastores report 1025
  - file system reports 1027
  - managing resources 266
  - NAS mounts report 1027
  - performance chart 1017, 1019
  - provisioning block storage 267
  - provisioning file storage 268
  - provisioning storage 80
  - SCSI Volumes (LUNs) report 1025
  - storage mapping to Tivoli Storage Productivity Center 1015
  - storage system metrics 1019
  - viewing block storage reports 1025
  - viewing fabric connections 1024
  - viewing file system reports 1026
  - viewing storage mapping 1014
  - viewing storage reports 80
  - viewing storage system metrics 1017
- vSphere Web Client extension reports 1013

## W

- web-based GUI
  - acknowledging alerts 184

- web-based GUI (*continued*)
  - acknowledging the status of a resource 202
  - acknowledging the status of a top-level resource 202
  - acknowledging the status of an internal resource 204
  - alert severities 161
  - alerts
    - acknowledging 184
    - removing 184
  - customizing lists 287
  - data collection
    - creating performance monitors 152
    - modifying performance monitors 154
    - performance monitors 151
    - starting performance monitors 156
    - stopping performance monitors 156
    - verifying that a performance monitor is running 153
    - verifying that a probe is running 148
  - display differences with the stand-alone GUI 195
  - exporting information about applications and departments 237
  - exporting information about resources, jobs, and alerts 285
  - fabrics 85
  - filtering 287
  - how information is organized 185
  - how the status of resources is determined 197
  - internal and related resources, list 129
  - managing tasks 289
  - Monitoring the status of resources 196
  - Opening the management GUI for a resource 284
  - removing alerts 184
  - resources that can be monitored 129
  - saving information about applications and departments 237
  - saving information about resources, jobs, and alerts 285
  - status of resources 196
  - switches 85
  - tasks
    - analysis execution 298, 299
    - creating provisioning tasks 294
    - creating tiering storage, balancing pools, and transforming-storage tasks 296
    - deleting 301
    - exporting provisioning tasks 295
    - implementing optimization recommendations 298
    - managing provisioning tasks 293
    - managing tiering storage, balancing pools, and transforming-storage tasks 295

- web-based GUI (*continued*)
  - tasks (*continued*)
    - pausing, resuming, and canceling analysis-execution tasks 300
    - running provisioning tasks 294
    - running tiering analysis, transform plan, and balance-analysis tasks 297
    - scheduling the implementation of optimization recommendations 299
  - viewing all performance monitors 157
  - viewing all tasks 290
  - viewing information about internal and related resources 194
  - viewing information about resources 185
  - viewing information about top-level resources 193
  - viewing pending tasks 290
  - viewing performance monitor logs 158
  - viewing performance monitors 157
  - viewing performance monitors for a resource 157
  - viewing task details 291
  - viewing task for a resource 291
  - viewing task for a resource type 290
  - viewing task logs 292
  - viewing tasks 290
  - viewing the overall status of resources 200
  - viewing the status of internal resources 201
  - viewing the status of top-level resources 201
  - zones 85
  - Windows 1137, 1139
  - wizards
    - Analyze Tiering 95
  - workload
    - volumes 276
  - WWN zoning 86

## Z

- zone 86
  - concepts 86
  - default 86
  - hardware 86
  - membership 86
  - non standard 88
  - port 86
  - software 86
  - supported switches 88
- zone set 86
  - temporary 86
- zones 85







Product Number: 5725-F92, 5725-F93, 5725-G33

Printed in USA

SC27-4060-05

