

IBM Spectrum Protect Snapshot
Version 8.1.7

*Installation and User's Guide for
Windows*



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 371.](#)

This edition applies to version 8, release 1, modification 7 of IBM Spectrum Protect Snapshot (product numbers 5725-X22, and 5608-AB8). It also applies to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2001, 2019.**

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

Contents

Figures.....	vii
Tables.....	ix
About this publication.....	xi
Who should read this guide.....	xi
Publications	xi
Reading syntax diagrams.....	xi
Windows updates V8.1.7.....	xv
Chapter 1. Overview.....	1
Volume Shadow Copy Service framework.....	2
Data protection in VSS environments.....	2
Data backup overview.....	4
Database backup types.....	4
Data backup methods.....	7
Policy management with IBM Spectrum Protect Snapshot.....	10
Data restore overview.....	17
VSS fast restore processing.....	17
VSS instant restore processing.....	17
VSS backups that are restored to alternate databases.....	18
Exchange mailbox restore operations.....	18
IBM System Storage requirements.....	20
IBM Spectrum Protect Snapshot with IBM SAN Volume Controller and IBM Storwize V7000.....	21
Failover clustering and AlwaysOn Availability.....	23
Availability database backup operations.....	24
Availability database restore operations.....	25
Automated IBM Spectrum Protect server failover for data recovery.....	25
Chapter 2. Planning.....	27
Storage capacity requirements.....	27
Best practices for IBM Spectrum Protect Snapshot with IBM XIV 11.6 Real-time Compression	28
Chapter 3. Installing and upgrading.....	29
Prerequisites.....	29
Installing IBM Spectrum Protect Snapshot.....	30
Silently installing IBM Spectrum Protect Snapshot for Windows.....	32
Options in silent installations.....	33
Creating and testing a silent installation package on a DVD or a file server.....	35
Batch files usage in silent installations.....	36
Silent installation error messages.....	36
Silently installing IBM Spectrum Protect Snapshot on Windows Server Core.....	36
Silently installing IBM Spectrum Protect Snapshot on Windows Server Core with the setup program.....	36
Silently installing IBM Spectrum Protect Snapshot on Windows Server Core with the Microsoft Installer program.....	37
Upgrading IBM Spectrum Protect Snapshot.....	37
IBM Spectrum Protect Snapshot migration.....	38

Managing migrated backups to a Database Availability Group node.....	38
Uninstalling IBM Spectrum Protect Snapshot.....	39

Chapter 4. Configuring..... 43

Specifying configuration parameters for IBM Spectrum Protect	43
Specifying IBM Spectrum Protect Snapshot configuration and options files in non-default locations..	46
Setting user preferences.....	46
Data Protection Properties.....	47
Configuring IBM Spectrum Protect Snapshot in a stand-alone configuration.....	58
Configuring an IBM Spectrum Protect Snapshot remote system in a stand-alone configuration....	60
Configuring IBM Spectrum Protect Snapshot to integrate with IBM Spectrum Protect.....	61
Configuring an IBM Spectrum Protect Snapshot remote system to integrate with IBM Spectrum Protect.....	65
Configuring IBM Spectrum Protect Snapshot to restore mailboxes from mounted Exchange Server database files.....	66
Configuring node definitions.....	67
Proxy node definitions for VSS backups.....	67
Configuring the system that runs offloaded backups.....	69
Configuring to protect SQL Server data.....	70
Configuring IBM Spectrum Protect Snapshot for SQL Server clustered environments.....	70
Configuring IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core.....	76
Configuring your system for mailbox restore operations.....	83
Configuring for mailbox restore operations (Exchange 2016 and later).....	84
Examples of distributed VSS backups in Microsoft Exchange Database Availability Groups and Microsoft SQL AlwaysOn Availability Groups.....	84
Examples of IBM SAN Volume Controller and IBM Storwize V7000 configuration scenarios.....	86

Chapter 5. Protecting your data..... 89

Starting Microsoft Management Console	89
Starting the IBM Spectrum Protect Snapshot command-line interface	89
Getting help for IBM Spectrum Protect Snapshot commands.....	90
Determining managed storage capacity.....	90
Protecting Microsoft Exchange Server data.....	91
Prerequisites.....	91
Backing up Exchange Server data by using VSS.....	99
Mounting Exchange Server backups.....	100
Deleting Exchange Server backups.....	101
Setting data restore options in Microsoft Management Console.....	102
Restoring an Exchange Server database.....	104
Restoring a Database Availability Group database backup	106
Restoring mailbox data.....	107
Restoring mailbox messages interactively with the Mailbox Restore Browser	112
Restoring mailboxes directly from Exchange Server database files.....	117
Restoring a deleted mailbox or items from a deleted mailbox.....	117
Restoring mailboxes on remote systems.....	118
Protecting SQL Server data.....	118
Prerequisites.....	119
Verifying the integrity of legacy databases by using the checksum option.....	123
Creating VSS backups of SQL Server databases.....	123
Creating legacy backups of SQL Server databases.....	125
Cloning an SQL Server database.....	128
Deleting SQL Server backups.....	130
Deactivating legacy backups of SQL Server databases.....	131
Setting single-user mode for restore operations.....	132
Setting data restore options.....	132
Restoring SQL Server data.....	135
Restoring an SQL Server database to an alternate instance.....	137

Restoring the master database.....	139
Restoring SQL databases with full-text catalogs and indexes.....	141
Protecting SQL Server data in a Windows Server Core environment.....	141
Protecting custom application and file system data.....	143
Prerequisites.....	144
Backing up custom application and file system data.....	145
Implementing custom application and file system backup scenarios.....	146
Mounting custom application and file system backups.....	147
Deleting custom application and file system backups.....	148
Restoring custom application and file system data.....	148
Implementing custom application and file system restore scenarios.....	150
Mounting VSS snapshots to remote servers.....	152
Enabling Windows PowerShell Remoting for Remote Management and Remote Mounting.....	155
Adding remote systems.....	157
Viewing, printing, and saving reports.....	158
Generating group reports.....	158
Chapter 6. Automating tasks.....	159
Preparing to use Windows PowerShell cmdlets.....	159
Cmdlets for Microsoft Management Console.....	160
Cmdlets for protecting Microsoft Exchange Server data.....	161
Cmdlets for protecting Microsoft SQL Server data.....	162
Cmdlets for protecting custom application and file system data.....	164
Automating Microsoft Exchange Server tasks.....	166
Automating Microsoft SQL Server tasks.....	167
Automating custom applications and file system tasks.....	168
Scheduling tasks.....	169
Chapter 7. Troubleshooting.....	171
Diagnosing problems.....	171
Error log files for IBM Spectrum Protect Snapshot components.....	171
Trace files for IBM Spectrum Protect Snapshot components.....	172
Diagnosing VSS issues.....	173
Resolving reproducible problems.....	173
Troubleshooting VSS backup and restore operations.....	174
Troubleshooting mailbox restore errors.....	176
Deleting mailbox history information.....	180
Troubleshooting configuration errors in a failover clustered environment.....	181
Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000.....	181
Resolving problems with IBM Support.....	183
Viewing trace and log files.....	183
Gathering trace and log files for remote systems.....	185
Gathering information about Exchange or SQL Server with VSS before you call IBM.....	186
Viewing system information.....	187
Sending support files by email.....	188
Online IBM support.....	188
Chapter 8. Reference.....	189
Support for Microsoft Exchange 2016 and later versions.....	189
Commands for Exchange Server.....	190
Backup command.....	191
Delete backup command.....	196
Help command.....	202
Mount backup command.....	203
Policy commands for IBM Spectrum Protect Snapshot for Exchange.....	208
Query Exchange command.....	210
Query FCM command.....	213

Query Managedcapacity command.....	217
Query TDP command.....	218
Restore command.....	220
Restorefiles command.....	226
Restoremmailbox command.....	230
Set command.....	248
Unmount backup command.....	256
Commands for SQL Server.....	259
Backup command.....	259
Delete backup command.....	268
Help command.....	272
Mount Backup command.....	273
Policy commands for IBM Spectrum Protect Snapshot for SQL.....	278
Query FCM command.....	280
Query Managedcapacity command.....	284
Query SQL command.....	285
Query TDP command.....	291
Restore command.....	292
Restorefiles command.....	302
Set command.....	306
Unmount Backup command.....	313
Commands for custom applications and file systems.....	316
Backup command.....	316
Delete backup command.....	320
Help command.....	324
Policy commands for IBM Spectrum Protect Snapshot.....	324
Mount backup command.....	329
Query component command.....	335
Query config command.....	338
Query backup command.....	341
Query managedcapacity command.....	346
Restore command.....	349
Unmount backup command.....	356
Update config command.....	360
VSS policy commands.....	365
Appendix A. Accessibility.....	369
Notices.....	371
Glossary.....	375
Index.....	377

Figures

- 1. Sample DAG environment..... 9
- 2. Example of how IBM Spectrum Protect Snapshot distributes VSS backups..... 85
- 3. Typical DAG configuration..... 95
- 4. Example of backups that are distributed across DAG members..... 95
- 5. Another example of backups distributed across DAG members..... 96

Tables

1. IBM Spectrum Protect Snapshot components in Windows environment.....	1
2. IBM Spectrum Protect Snapshot for Exchange Server VSS backup types.....	5
3. IBM Spectrum Protect Snapshot for SQL Server backup types.....	6
4. Exchange Server Recoverable Items folder contents.....	20
5. Silent installation options.....	33
6. Silent installation features (base client only).....	34
7. Silent installation features (IBM Spectrum Protect Snapshot for Windows).....	35
8. Commands for creating a silent installation package.....	35
9. Diagnostics modes and their usage.....	49
10. Field entry in the IBM Spectrum Protect Node Names page.....	62
11. Field entry in the IBM Spectrum Protect Node Names page.....	63
12. Required node names for basic VSS operations.....	67
13. Required node names for basic VSS offloaded backups.....	68
14. Configuration options for file system and custom applications.....	72
15. IBM Spectrum Protect Snapshot help commands.....	90
16. Options for integrity checking.....	99
17. Database restore options.....	102
18. Database restore selection options.....	105
19. Database restore options.....	109
20. Restore options.....	109
21. Selecting mailboxes to restore.....	114
22. Previewing and filtering mailbox items.....	115
23. Restoring a mailbox to another mailbox or .pst file.....	116

24. Database backup views.....	124
25. Database backup views.....	126
26. Database backup options.....	126
27. Database restore options.....	132
28. Database backup views.....	136
29. Database restore selection options.....	136
30. Database backup views.....	138
31. Database backup views.....	138
32. Backup options.....	145
33. Database restore selection options.....	149
34. Cmdlets to protect Microsoft Exchange Server data.....	161
35. Cmdlets to protect Microsoft SQL Server data.....	162
36. Cmdlets to protect custom application and file system data.....	164
37. Mailbox restore options	189
38. MAPI/CDO changes	190
39. SQL Server connection protocols.....	266
40. SQL Server connection protocols.....	288
41. SQL Server connection protocols.....	300

About this publication

Use the IBM Spectrum Protect Snapshot software to create and manage volume-level snapshots while the applications that contain data on those volumes remain online.

Throughout this document, the term *Windows VSS System Provider* (unless otherwise specified) refers to the standard Windows System provider.

Since the previous edition, changes are marked with a vertical bar (|) in the left margin.

Who should read this guide

This publication is intended for administrators who are responsible for implementing a backup solution in database server environments.

It is assumed that you understand the following storage systems, operating systems, or applications, as applicable:

- The storage system that is used for the database:
 - Any storage devices that implement the VSS provider interface.
 - IBM® System Storage® DS3000, DS4000®, DS5000
 - IBM System Storage SAN Volume Controller
 - IBM Storwize® V7000
 - IBM XIV® Storage System
 - IBM System Storage DS8000® series
- Windows operating system
- Microsoft Volume Shadow Copy Service (VSS)
- Microsoft Exchange Server
- Microsoft SQL Server
- Active Directory

Publications

The IBM Spectrum Protect product family includes IBM Spectrum Protect Plus, IBM Spectrum Protect for Virtual Environments, IBM Spectrum Protect for Databases, and several other storage management products from IBM.

To view IBM product documentation, see [IBM Knowledge Center](#).

Reading syntax diagrams

To read a syntax diagram for entering a command, follow the path of the line. Read from left to right and from top to bottom.

- The ►— symbol indicates the beginning of a syntax diagram.
- The —► symbol at the end of a line indicates the syntax diagram continues on the next line.
- The ►— symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The —►◀ symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or variable, can be:


- On the line (required element)
- Above the line (default element)
- Below the line (optional element)

Syntax diagram description	Example
<p>Abbreviations:</p> <p>Uppercase letters denote the shortest acceptable truncation. If an item displays entirely in uppercase letters, it cannot be truncated.</p> <p>You can type the item in any combination of uppercase or lowercase letters.</p> <p>In this example, you can enter KEYWO, KEYWORD, or KEYWOrd.</p>	<p>►► KEYWOrd ◄◄</p>
<p>Symbols:</p> <p>Enter these symbols exactly as they display in the syntax diagram.</p>	<p>* Asterisk</p> <p>{ } Braces</p> <p>:</p> <p>Colon</p> <p>,</p> <p>Comma</p> <p>= Equal Sign</p> <p>- Hyphen</p> <p>0 Parentheses</p> <p>.</p> <p>Period</p> <p>' Single quotation mark</p> <p>Space</p> <p>" Quotation mark</p>
<p>Variables:</p> <p>Italicized lowercase items (<i>var_name</i>) denote variables.</p> <p>In this example, you can specify a <i>var_name</i> when you enter the KEYWORD command.</p>	<p>►► KEYWOrd — <i>var_name</i> ◄◄</p>


Syntax diagram description**Example**

Repetition:


An arrow that returns to the left means you can repeat the item.



A character or space within an arrow means you must separate the repeated items with that character or space.

**Required Choices:**


When two or more items are in a stack and one of them is on the line, you must specify one item.




In this example, you must choose A, B, or C.

Optional Choice:


When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all.



When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.

**Defaults:**


Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line.



In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.

Repeatable Choices:



A stack of items followed by an arrow that returns to the left means you can select more than one item. In some cases, it means you can repeat a single item.



In this example, you can choose any combination of A, B, or C.

Syntax Fragments:

Some diagrams because of their length, must fragment the syntax. The fragment name is displayed between vertical bars in the diagram. The expanded fragment is displayed between vertical bars in the diagram after a heading with the same fragment name.


**The fragment name**

Syntax diagram description**Example**

Footnote:

A footnote in the diagram references specific details about the syntax that contains the footnote.

In this example, the footnote by the arrow references the number of times you can repeat the item.



Notes:

¹ Specify *repeat* as many as 5 times.

IBM Spectrum Protect Snapshot for Windows updates V8.1.7

This document provides information about what's new or what has changed in IBM Spectrum Protect Snapshot version 8.1.7.

New and changed information is indicated by a vertical bar (|) to the left of the change.

What's new

Microsoft Windows Server 2019

With IBM Spectrum Protect Snapshot Version 8.1.7, Windows Server 2019 is supported.

Microsoft Exchange Server 2019

With IBM Spectrum Protect Snapshot Version 8.1.7, Exchange Server 2019 is supported.

Mailbox Restore Browser support for Exchange 2016 and Exchange 2019

Browse and restore individual mailbox items by using the Mailbox Restore Browser. This feature is now supported on Exchange Server 2016 and Exchange Server 2019. For more information, see [“Configuring your system for mailbox restore operations \(Exchange 2016 and later\)” on page 84](#)

Chapter 1. Overview

With IBM Spectrum Protect Snapshot, you can back up and restore Exchange Server data, SQL Server data, file system volumes, and custom application data on volumes. You can back up and restore the data to IBM Spectrum Protect server or local shadow volumes.

A *local shadow volume* contains data that is stored on shadow volumes, which are local to a disk storage system.

You can install, configure, and use IBM Spectrum Protect Snapshot with different operating systems, databases, and applications or file systems. The software is compatible with most hardware that uses FlashCopy® technology.

The following table lists IBM Spectrum Protect Snapshot components that operate in Microsoft Exchange Server Microsoft SQL Server and file system and custom application environments:

Component	Description
Microsoft Management Console (MMC) Snap-in and Base System Services	Uses MMC and a generic backup agent to create snapshots of file systems, applications, or databases.
VSS Requestor	Uses the VSS backup-archive client as a VSS interface to communicate with Microsoft VSS services, access data, and create volume shadow copies.
IBM Spectrum Protect Snapshot for Exchange Server	Uses Microsoft Exchange Server capabilities to complete the following tasks: <ul style="list-style-type: none">• Store VSS backup copies locally as persistent snapshots, or store VSS backup copies as snapshots on the IBM Spectrum Protect server. You can also offload backups to IBM Spectrum Protect.• Centralize policy management and scheduling.• Complete VSS, and volume-level instant restore operations.• Complete individual mailbox restore operations.
IBM Spectrum Protect Snapshot for SQL Server	Uses Microsoft SQL Server capabilities to complete the following tasks: <ul style="list-style-type: none">• Complete legacy-style and VSS backups.• Store VSS backup copies locally as persistent snapshots, or store VSS backup copies as snapshots on the IBM Spectrum Protect server. You can also offload backups to IBM Spectrum Protect.• Centralize policy management and scheduling.• Complete VSS, and volume-level instant restores.

You can use IBM Spectrum Protect Snapshot with IBM Spectrum Protect. When the two products are used together, they provide advanced data protection and centrally managed, policy-based administration capabilities for IBM Spectrum Protect Snapshot backup images.

Volume Shadow Copy Service framework

Volume Shadow Copy Service (VSS) provides a common interface model to generate and manage online snapshots of Exchange Server, SQL Server, file system, or custom application data.

The Microsoft VSS service manages and directs three VSS software components that are used during VSS operations: the VSS writer, the VSS Requestor, and the VSS provider. The VSS writer is the application that stores data on the source volumes. The VSS Requestor is the backup software. The VSS provider is the combined hardware and software that generates the snapshot volume.

The VSS system provider creates and maintains snapshots on local shadow volumes and refers to the default VSS provider that is available with Windows Server. If you use the Windows VSS system provider, no configuration is required. However, you can make changes by using the **VSSADMIN** commands.

With a VSS hardware and software copy provider, you can create shadow copies of running volumes on demand. A hardware provider uses a hardware storage adapter or controller to manage shadow copies at the hardware level. IBM Spectrum Protect Snapshot for Windows software does not control the VSS hardware provider. The VSS hardware provider is controlled by the hardware vendor. Install and configure the VSS hardware and software provider as required. When you use a hardware provider, configure the disks as basic disks. Do not use dynamic disks.

Data protection in VSS environments

The characteristics of Volume Shadow Copy Service (VSS) backup and restore operations can affect management tasks, for example, the backup types that you can run, the backup granularity, and the backup storage location options.

As you decide your backup and restore strategies, be aware of VSS requirements and guidelines.

VSS backup characteristics

Backups can be stored on local shadow volumes, on an IBM Spectrum Protect server, or at both locations. You can define different policy settings for each backup location.

Databases must have unique names. If a database has the same name as another database, but the capitalization differs, the software does not differentiate between case.

When you back up Exchange Server data by using IBM Spectrum Protect Snapshot, VSS backups have the following characteristics:

- Backups provide an Exchange Server database integrity check function, but do not provide a zeroing function.
- You can run full, copy, differential and incremental backups.
- You can restore a backup to a local disk only on the same system.

When you back up SQL Server data by using IBM Spectrum Protect Snapshot, VSS backups have the following characteristics:

- Backups can run in a valid Microsoft Windows Failover Clustering or Veritas Cluster Server (VCS) environment.
- You can only run full and copy-only full backups.
- You can run legacy differential and legacy log backups after you restore a full VSS backup.
- You can back up SQL server databases in an AlwaysOn Availability Group (AAG).
- To relieve resources on production servers, you can offload backups from IBM Spectrum Protect server storage to another system.

For custom application and file system data, data is backed up at the file system level. You can use drives and mount points. You can back up data in a valid Microsoft Windows Failover Clustering or Veritas Cluster Server (VCS) environment.

Related concepts

Failover clustering and AlwaysOn Availability

In an SQL Server cluster environment, two different kinds of clustering are supported; AlwaysOn Failover Cluster Instances (FCI) and AlwaysOn Availability Groups (AAGs). An SQL AlwaysOn failover cluster instance provides high availability and disaster recovery at the SQL Server level. AlwaysOn Availability Groups (AAG) provide high availability and disaster recovery at SQL database level. IBM Spectrum Protect Snapshot for SQL Server protects availability databases in both AlwaysOn failover cluster instances and in an AAG.

VSS backup requirements

You can plan your VSS backup strategy to optimize the performance of your backup operations and to avoid potential problems. Follow these guidelines when you plan your VSS backups:

• Planning VSS backups

- Use custom-application VSS backups and files-system data VSS backups for only NTFS and ReFS volumes.

For file system and custom application data, data is restored at the volume level. When a file system or custom application data is restored, all files from the VSS snapshot backup are restored to their original location.

- Use single hardware Logical Unit Numbers (LUN) for log files, system files, and database files.
- Use basic disks, which are initialized for basic storage. A basic disk consists of basic volumes, such as primary partitions, extended partitions, and logical drives. Do not use dynamic disks.
- If you plan to keep VSS snapshot backups only on local shadow volumes, know how to implement the configuration options of your VSS hardware provider.

For example, if your VSS hardware provider supports a full-copy snapshot versus a copy-on-write snapshot mechanism, full-copy type implementations have greater disk storage requirements. However, full-copy type implementations do not rely on the original volume to restore the data and are less risky. Copy-on-write implementations require less disk storage but rely on the original volume to restore the data.

- Do not place multiple volumes on the same LUN. Configure a single volume, single partition, and single LUN as one-to-one.
- Do not set the ASNODENAME option in the `dsm.opt` file when you use IBM Spectrum Protect Snapshot. Setting the ASNODENAME option can cause VSS data backups and VSS restore operations to fail.

• Running parallel VSS backups

If you need to run parallel VSS backups, do the following:

- Stagger the start time of the backups by at least 10 minutes. This interval ensures that the snapshot operations do not overlap.



Attention:

If backup operations overlap, a VSS timeout error may occur and the second backup request may fail. Therefore, it is recommended to stagger the start time of the backups.

- Configure the parallel instance backups so that snapshots of the same volumes are not created.
- Ensure that parallel backups do not create a snapshot of the same LUN.

VSS restore characteristics

In a VSS restore operation, VSS backups (Exchange or SQL Server database files and log files) on IBM Spectrum Protect server storage are restored to their original location on the Exchange Server or SQL Server. For SQL Server VSS backups, you can also restore to an alternate SQL Server instance.

VSS restore requirements

Unless otherwise specified, a *VSS restore* operation refers to all restore types that use VSS, including VSS restore, VSS fast restore, and VSS instant restore operations.

If you complete VSS snapshot backups with the backup destination parameter set to TSM, *restore* processing also refers to an image-level restore from the IBM Spectrum Protect server.

As you decide your restore strategies, be aware of VSS requirements.

VSS instant restore operations

A VSS instant restore operation overwrites the entire contents of the source volumes.

- If you do not want to overwrite the source volumes, ensure that you set the **Instant Restore** option to **No** in Microsoft Management Console (MMC).
- VSS instant restore processing requires that the local disk is not accessed by other applications, for example, Windows Explorer.
- When you run a VSS instant restore operation, verify that there is no other data on the volumes that are being restored.
- Before you start a VSS instant restore operation, ensure that any previous background copies that contain the volumes that are being restored are completed. XIV, SAN Volume Controller, or Storwize family with space-efficient target volumes do not need to be completed.
- To restore VSS snapshots on non-IBM storage systems, you must create the snapshot as a transportable snapshot and import the transportable snapshot into the system.

Tip: For instant restore operations on non-IBM storage devices, the VSS Hardware provider must enable the device to implement and support the Microsoft VSS ResyncLuns API. In addition, the VSS Hardware provider must support creating transportable snapshots during backups, and importing transportable snapshots during restores. You cannot perform instant restore operations on these devices unless the VSS ResyncLuns API is implemented.

Tip: For instant restore operations on IBM storage devices, you can use either private interfaces or the VSS ResyncLuns API. The format that is used depends on the device. When you use private interfaces on these devices, it is not necessary to set the shadow copy as a transportable snapshot. When you use VSS ResyncLuns API on these devices, it is necessary to set the shadow copy as a transportable snapshot.

VSS fast restore operations

In a VSS fast restore operation, if you do not want to overwrite all the files on the original volume, mount the snapshot. Copy only the files that you want to restore.

When you restore data, only use basic disks that are initialized for basic storage. A basic disk consists of basic volumes, such as primary partitions, extended partitions, and logical drives. Do not use dynamic disks.

When you complete a VSS restore operation from local shadow volumes, the bytes that transfer are displayed as 0 because no data (0) is restored from the IBM Spectrum Protect server.

Data backup overview

IBM Spectrum Protect Snapshot can use the Microsoft Volume Shadow Copy Service (VSS) framework to produce a point-in-time, consistent, online copy of Exchange Server, SQL Server, file system, or custom application data.

Database backup types

With IBM Spectrum Protect Snapshot for Exchange Server and IBM Spectrum Protect Snapshot for SQL Server, you can use the common interface in the Volume Shadow Copy Service (VSS) framework to create database backups.

You can back up IBM Spectrum Protect Snapshot for Microsoft Exchange Server data by using the following methods:

Table 2. IBM Spectrum Protect Snapshot for Exchange Server VSS backup types

VSS backup types	
Full backup	With this method, IBM Spectrum Protect Snapshot for Microsoft Exchange Server backs up the specified database and associated transaction logs. If the database is not mounted, the backup fails and the transaction logs are not truncated.
Copy backup	With this method, IBM Spectrum Protect Snapshot for Microsoft Exchange Server does not delete transaction log files after the backup. Otherwise, this type is similar to a full backup. Use a copy backup to create a full backup of the Exchange Server database without disrupting any backup processes that use an incremental or differential backup.
Incremental backup	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft Exchange Server backs up only transaction logs. Transaction log files are not deleted if the backup fails.</p> <p>When you restore an Exchange Server database from an incremental backup, you must complete the following tasks:</p> <ul style="list-style-type: none"> • Restore the last full backup. • Restore any other incremental backups that occur between the full backup and the incremental backup. • Restore the incremental backup.
Differential backup	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft Exchange Server backs up transaction logs.</p> <p>When you follow a full backup with only differential backups, the last full backup and the last differential backup contain all the data that is required to restore the database to the most recent state.</p> <p>When you restore an Exchange Server database from a differential backup, you must complete the following tasks:</p> <ul style="list-style-type: none"> • Restore the last full backup. • Restore this differential backup, but no other differential backups.



Attention: When you enable circular logging, you cannot use differential or incremental backups. Data loss might occur if the log wraps before the incremental or differential backup ends.

VSS backups are at the volume and file-level. Legacy backups are a stream of bytes that IBM Spectrum Protect Snapshot for Microsoft SQL Server stores on the IBM Spectrum Protect server.

You can back up IBM Spectrum Protect Snapshot for Microsoft SQL Server data by using the following methods:

Table 3. IBM Spectrum Protect Snapshot for SQL Server backup types

IBM Spectrum Protect Snapshot for SQL Server	
Full database backup (Legacy and VSS)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server backs up an SQL Server database and the portion of the transaction log that is necessary to provide a consistent database state. With this backup type, the copy includes enough information from any associated transaction log to create a backup that is consistent with itself. The portion of the log that is included contains only the transactions that occur from the beginning of the backup until its completion.</p>
Copy-only full backup (Legacy and VSS)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server creates data backups that do not affect existing backup and restore processes and can be retained in the longer term. For example, you can use this type to back up a log before an online file restore operation. In this example, the copy-only full backup is used once. After the backup is restored, it is deleted.</p>
Differential backup (only Legacy)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server backs up only the data pages in an SQL Server database instance that changed after the last full backup. A portion of the transaction log is also backed up.</p> <p>Differential backup is associated with the last full backup that was run. The last full backup might be completed by IBM Spectrum Protect Snapshot for Microsoft SQL Server or another application. For example, if you run a full SQL Server-to-disk backup, and run a differential backup, the differential backup is associated with the SQL Server disk backup.</p> <p>You cannot use differential backup for databases on the secondary replica in Microsoft SQL Server.</p>

Table 3. IBM Spectrum Protect Snapshot for SQL Server backup types (continued)

IBM Spectrum Protect Snapshot for SQL Server	
Log backup (only Legacy)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server backs up only the contents of an SQL Server database transaction log since the last successful log backup. This type of backup is preceded by a full backup or an equivalent type of backup.</p> <p>Log backups normally follow full backups. The portion of the log that is included in full and differential backups is not equivalent to a log backup. Additionally, in full and differential backups, the log is not truncated as it is during a log backup. However, a log backup that follows a full or differential backup includes the same transactions as a full or differential backup. Log backups are not cumulative as are differential; they must be applied against a base backup and in the correct order.</p>
File backup (only Legacy)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server backs up only the contents of a specified SQL Server logical file. This type of backup can ease the scheduling conflicts if you must back up large databases. You can back up different sets of files during different scheduled backups. File, group, and set backups must be followed by a log backup, but a full backup is not required.</p>
Group backup (only Legacy)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server backs up only the contents of a specified SQL Server file group. You can back up the set of database tables and indexes within a specific group of files.</p> <p>The group is specified as part of the setup within SQL Server when you define the database files. If no group is specified and all the database files are part of the primary group, you cannot partially back up or partially restore the database by using the group.</p>
Set backup (only Legacy)	<p>With this method, IBM Spectrum Protect Snapshot for Microsoft SQL Server backs up the contents of specified SQL Server file groups and files as a unit.</p>

Data backup methods

With IBM Spectrum Protect Snapshot, you can use Volume Shadow Copy Service (VSS) to back up Exchange Server and SQL Server data. For SQL Server, you can also run legacy backups that create a copy of all or part of an SQL database or logs on IBM Spectrum Protect storage media.

You can run Exchange Server backup operations in a Database Availability Group (DAG) environment.

VSS data backups

You can store VSS backups on local VSS shadow volumes, or, when integrated with IBM Spectrum Protect, in IBM Spectrum Protect server storage.

VSS backups eliminate the need for the server or file system to be in backup mode for an extended time. The length of time to complete the snapshot is measured in seconds, not hours. In addition, a VSS backup allows a snapshot of large amounts of data at one time because the snapshot works at the volume level.

You must ensure that sufficient space is available for the snapshot at the storage destination. Both storage destinations require space to store the snapshot until the data transfer to the IBM Spectrum Protect server is complete. After the data transfer to the server is complete, VSS backups that are stored locally on VSS shadow volumes are directly accessible by the system. The snapshot volume is released and the space can be reused.

- For data that is backed up to local VSS shadow volumes, the snapshot backup is on the shadow copy volume.
- For data that is backed up only to IBM Spectrum Protect server storage, a local snapshot backup is run and the data on the local snapshot volume is sent to the IBM Spectrum Protect server.
- For data that is backed up to VSS shadow volumes and IBM Spectrum Protect server, the local snapshot volume is retained as a local backup after the transfer to the IBM Spectrum Protect server is complete.

If you store VSS backups both locally and to IBM Spectrum Protect server, and the maximum number of local backup versions to be maintained is reached, the oldest local backup version expires to create the new snapshot for the backup to IBM Spectrum Protect server storage. The maximum number of local backup versions that are maintained is set in the IBM Spectrum Protect policy.

Offloaded VSS backups

By running an offloaded backup, you can move the backup load from the production system to another system. You can reduce the load on network, I/O, and processor resources during backup processing.

Use the **RemoteDSMAGENTNode** parameter to run an offloaded system. Ensure that you install a VSS hardware provider, which supports transportable shadow copy volumes, on the production and secondary systems.

Exchange Database Availability Group backups

You can use the high-availability feature of Database Availability Group (DAG) backups for enhanced data and service availability, and automatic recovery from failures. You can use Exchange Server 2013 or later versions with DAG backups to improve Exchange Server data backups and data recovery.

Beginning with Exchange 2013 SP1, you can also back up Exchange Server databases in a Database Availability Group (DAG) environment without a Cluster Administrative Access Point (CAAP).

A DAG environment includes the following functions:

- A group of up to 16 mailbox servers that can host up to 100 mailbox databases
- Up to 16 online copies of a database (1 active database and up to 15 passive databases)
- Synchronous or lagged replication. With lagged replication, you can delay the replaying of logs on target databases if, for example, there are time differences between source and target databases.
- Automatic migration and failover of active database copies

The following figure illustrates a DAG environment:

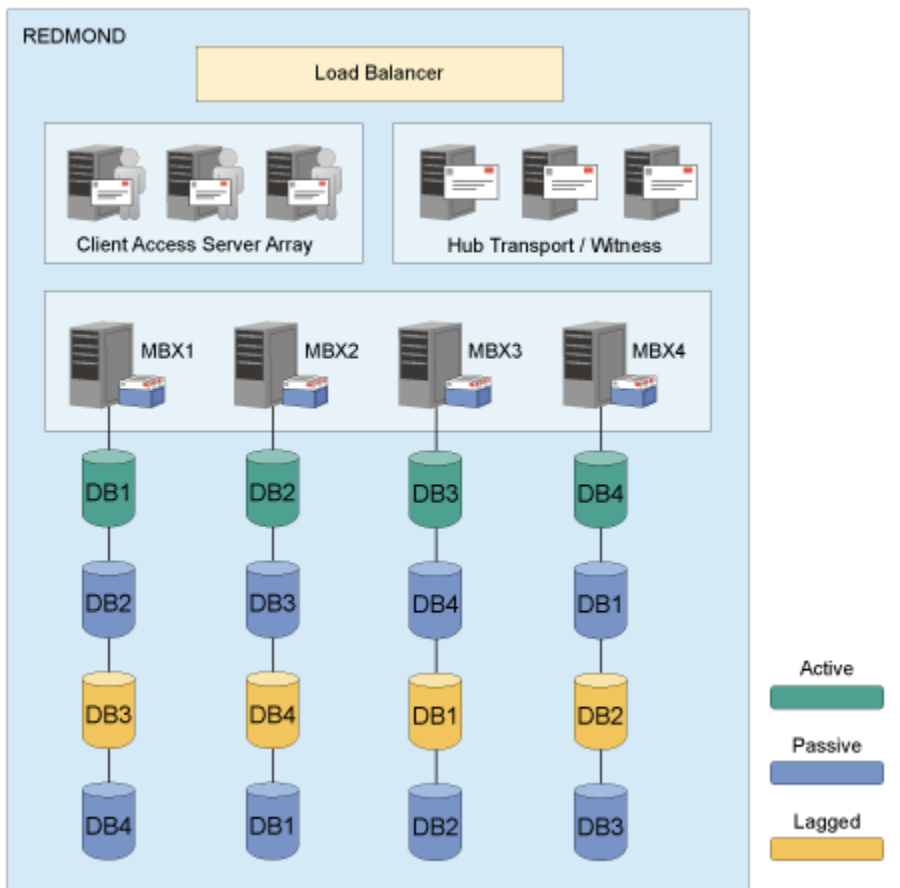


Figure 1. Sample DAG environment

Database copies are mirrored on any node within the DAG. You can complete the following tasks:

- Query DAG database copies, including status.
- Manage full, copy, incremental, and differential backups of active and passive databases within a DAG. You can create a backup from any active database copy, any passive synchronous copy, or any lagged copy within the DAG. If you back up a lagged database copy, it might take more time to restore the backup because the lagged copy can have more transaction logs to restore and replay. As a best practice, create your backup from a passive synchronous copy and not a lagged copy.
- Move an active database copy to other nodes.
- Query all DAG database copy backups.
- Restore all DAG database copy backups.
- Restore data into an active database, from either active or passive database copy backups.
- Restore data into a recovery or alternate database.
- Process Individual Mailbox Restore (IMR) operations from a DAG database copy backup.
- Delete DAG database copy backups.

SQL Server legacy backups

With IBM Spectrum Protect Snapshot for SQL Server, you can run legacy backups and store the backup on IBM Spectrum Protect server.

Legacy backups are unlike VSS backups because volume and file-level data are not backed up with this method.

Policy management with IBM Spectrum Protect Snapshot

With IBM Spectrum Protect Snapshot, you can manage and configure storage management policies for backups. A backup policy determines how backups on local shadow volumes are managed and retained.

Policy definitions apply only when you use a stand-alone configuration. If you configure IBM Spectrum Protect Snapshot to use the IBM Spectrum Protect server, the policy definitions are defined on the IBM Spectrum Protect server. VSS policy bindings are still managed locally.

IBM Spectrum Protect Snapshot uses a policy to determine how backups are retained. With IBM Spectrum Protect Snapshot, you can create, change, and view policies, and set binding policy statements to manage your backups.

Although IBM Spectrum Protect policy determines how IBM Spectrum Protect Snapshot backups are managed on IBM Spectrum Protect storage, backup retention on local shadow volumes is determined by version and time-based policies. Ensure that sufficient local storage space is available on local shadow volumes for a VSS backup. In addition, verify that enough available storage space is assigned to the volumes to accommodate your backup operations. The shadow copy volume that is the storage destination of a snapshot must have sufficient space for the snapshot.

Environment and storage resources also affect how many backup versions are maintained on local shadow volumes. The amount of space that is required depends on the VSS provider that you use.

How backups expire based on policy

Backups expire based on IBM Spectrum Protect Snapshot policy.

Expiration is the process by which SQL Server, Exchange Server, or custom application and file system backup objects are identified for deletion when the expiration date is past or the maximum number of backup versions that must be retained is reached.

The date on which data expires depends on the business needs that are identified by the recovery point objective (RPO) and the recovery time objective (RTO) of your enterprise. For example, legal, operational, and application requirements affect how data must be protected to meet these RPO and RTO demands. With IBM Spectrum Protect Snapshot, you can specify the number of snapshot backups to retain and the length of time to retain them.

Backups can expire during a query, backup, or restore operation of an IBM Spectrum Protect Snapshot session.

You specify the number of backup copies that are retained. When the maximum number of backup copies is reached, the oldest backup expires and is deleted. You can specify the maximum number of backup copies in an IBM Spectrum Protect Snapshot policy.

A backup copy is retained for a maximum number of days. The maximum number of days that a backup can be retained is specified in the IBM Spectrum Protect Snapshot policy.

How policy affects backup management on IBM Spectrum Protect Snapshot for Windows

An IBM Spectrum Protect policy determines how IBM Spectrum Protect Snapshot for Microsoft Exchange Server and IBM Spectrum Protect Snapshot for Microsoft SQL Server backups are managed on IBM Spectrum Protect storage and on local shadow volumes when the environment is configured for VSS operations.

The IBM Spectrum Protect server recognizes IBM Spectrum Protect Snapshot for Microsoft Exchange Server and IBM Spectrum Protect Snapshot for Microsoft SQL Server as a *node*.

Data that is backed up to IBM Spectrum Protect storage from the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server node is stored and managed according to settings that you specify in the IBM Spectrum Protect server policy.

The IBM Spectrum Protect policy manages the VSS backups that are placed in IBM Spectrum Protect server storage pools. The IBM Spectrum Protect server manages VSS backups.

If you use IBM Spectrum Protect for Copy Services and upgrade to IBM Spectrum Protect Snapshot for Exchange Server or IBM Spectrum Protect Snapshot for SQL Server, with the license for IBM Spectrum Protect for Copy Services, you can store VSS backups to local shadow volumes.

IBM Spectrum Protect requires that sufficient storage space is available to create shadow volumes for VSS backup processing. Even when the VSS backup destination is the IBM Spectrum Protect server, storage space to create a shadow volume is still required temporarily.

The number of local backup versions that are maintained by the IBM Spectrum Protect server is determined by the value that is specified by the IBM Spectrum Protect server **verexists** parameter, which is defined in the copy group of the management class to which the local backup belongs. It is not necessary to allocate target sets when you use the VSS system provider. When you do not use the VSS system provider, the number of target volume sets that are allocated for local backups must be equal to the value of the **verexists** parameter. Target volume sets are not applicable to IBM XIV Storage Systems.

For example, if `verexists=3`, then at least three sets of target volumes must be allocated for the backup to complete successfully. If only two sets of target volumes are allocated, the third and subsequent backup attempts fail. If more sets of target volumes exist than the number specified by the **verexists** parameter, these sets are ignored by the IBM Spectrum Protect server. A high number of local backup versions cannot be stored. If you want to have n number of local backup versions, set the **verexists** parameter to $n + 1$.

If you keep only one backup, the same disk is reused. The process initially removes the existing backup and attempts the new backup. If the new backup fails, no backups exist.

If you retain multiple backups, the oldest backup is removed before another backup is created. If the new backup fails, you might have one less backup than specified by the policy. For example, if you specify that you want to retain five backups, but the last backup fails, you might have only four backup versions.

Ensure that you specify a **verexists** value that meets your VSS backup goals. If you have limited storage space for VSS operations and are restricted to a `verexists=1` setting, set the backup destination to BOTH. This option stores the backup on local shadow volumes and sends a copy to IBM Spectrum Protect server storage.

You can change and delete VSS backups that IBM Spectrum Protect Snapshot for Microsoft Exchange Server and IBM Spectrum Protect Snapshot for Microsoft SQL Server creates and stores on local shadow volumes. From the command-line interface, for example, issue the Microsoft **VSSADMIN DELETE SHADOWS** command to remove a VSS backup that is managed by IBM Spectrum Protect.

IBM Spectrum Protect is not able to prevent the backup from being removed. In this instance, IBM Spectrum Protect detects that the backup is removed and reconciles its index of available backups with what is on local shadow volumes. Because backups can be removed, establish a strategy that protects VSS backup data that is stored on local shadow volumes from being compromised.

When you use the configuration wizard in the GUI, the **VSSPOLICY** parameter is set in the `tdpexc.cfg` or `tdpsql.cfg` file.

Restriction:

If you configure IBM Spectrum Protect Snapshot to integrate with IBM Spectrum Protect server, do not simultaneously configure the following items:

- In the VSS Requestor options file (`baclient\dsm.opt`), do not specify the following entry:

```
Include.Image volume management-class-name
```

- In the IBM Spectrum Protect Snapshot configuration file (`fcmcfg.xml`), Exchange configuration file (`tdpexc.cfg`), or SQL configuration file (`tdpsql.cfg`), do not specify VSSPOLICY statements that use the TSM option to back up data to IBM Spectrum Protect.

Depending on the policy management settings, you can reuse a logical unit number (LUN) for a new backup. When a backup is requested and the maximum number of versions is reached, the software deletes the oldest snapshot (backup) to make space for the snapshot. If the new request fails after the oldest snapshot is deleted, you have one less backup version than expected.

You must manage the policy for local backups to reconcile the local backup repository with the information that is stored on the IBM Spectrum Protect server. For example, if target volume LUNs that are used for a local backup are removed from the storage system, the information that represents the backup on the IBM Spectrum Protect server must be reconciled. Similarly, if an IBM Spectrum Protect

server policy determines that a local backup copy is no longer needed, the local backup manager must free the target volume LUNs to the storage system. The local backup manager is released so that these LUNs can be used for future backup operations. IBM Spectrum Protect automatically detects when these situations occur and completes the reconciliation.

Consider the scenario where you use a two-member Database Availability Group (DAG), named MEMBER1 and MEMBER2. When you complete a backup to LOCAL on MEMBER1 and complete more backups on MEMBER2, the backups to LOCAL on MEMBER1 do not expire until the next time you back up, query, or delete data on MEMBER1. In this scenario, you might use more storage than specified by **verexists**.

Creating a local backup policy

A local backup policy determines how different backup versions are retained on local shadow volumes.

Before you begin

Backup retention on local shadow volumes is determined by your overall backup strategy, the type and number of VSS backup version on IBM Spectrum Protect and on the local shadow volumes, and time-based policies. Ensure that there is sufficient local storage space on local shadow volumes. The amount of space that is required depends on the VSS provider that you use.

About this task

When IBM Spectrum Protect Snapshot is connected to an IBM Spectrum Protect server, the backup policy is defined by the server. When IBM Spectrum Protect Snapshot is configured in stand-alone mode, you can define the backup policy.

Use the following steps to create and manage local backup policies.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the navigation tree, click **IBM Spectrum Protect**.
3. Select an **Exchange Server, SQL Server, or File System** instance.
4. In the **Actions** pane, click **Properties**.
5. From the list of available property pages, select **Policy Management**.
6. Add, delete, or update local policies for data retention.

When you add a policy, specify a unique policy name. Double-click the policy to edit a policy field. To retain an unlimited number of snapshots, or to retain snapshots for an unlimited number of days, specify NL.

7. Click **Save**.

What to do next

After you add a policy, you can bind a backup to that policy. Updates to existing, bound policies do not take effect until the next backup is run.

Specifying policy binding statements

Bind policy statements to associate Microsoft SQL Server, Microsoft Exchange Server, and custom application and file system backups to a management policy.

About this task

A default policy binds any backups that are not explicitly bound to a named policy.

For custom application and file system backups, policy-binding statements are stored in the IBM Spectrum Protect Snapshot configuration file, `fcmcfg.xml`, by default.

For Exchange Database Availability Groups (DAG), all the DAG members that share the DAG node must use the same VSS policy.

Tip: Use the same policy binding method for SQL Server, Exchange Server, or custom application and file system backups. Define a policy statement in the respective configuration file. A custom application or file system statement identifies the name of the volume or mount point directory (component) instead of the name of the database (object name).

Procedure

1. Specify the policy-binding statements to use to bind snapshots to a policy. Manually add the binding statements in the respective configuration file that defines the policy statements.

Policy-binding statements in the IBM Spectrum Protect Snapshot for Exchange Server or IBM Spectrum Protect Snapshot for SQL Server configuration files might look similar to the information in the following table.

	<i>server name</i>	<i>object name</i>	<i>backup type</i>	<i>backup dest</i>	<i>mgmt class</i>
VSSPOLICY	*	"Accounting"	FULL	LOCAL	MC_1
VSSPOLICY	SERVER_3	"Human Resources"	INCR	LOCAL	MC_6

2. For custom application and file system backups, modify the default `fcmcfg.xml` configuration file by issuing the following commands:

- Enter **FCMCLI INSERT VSSPOLICY** to insert a VSS policy binding statement at the position that is specified by the **SEQnumber** parameter.
- Enter **FCMCLI UPDATE VSSPOLICY** to update an existing VSS policy binding statement at the position that is specified by the **SEQnumber** parameter.
- Enter **FCMCLI DELETE VSSPOLICY** to delete an existing VSS policy binding statement at the position that is specified by the **SEQnumber** parameter.

For custom application and file system data, the following sample command inserts a new VSS policy binding statement at the position that is specified by the **SEQnumber** parameter:

```
FCMCLI INSERT VSSPOLICY "*" L:\mountdir FULL LOCAL MC1Q11" /SEQnumber=2
```

where:

- the asterisk character (*) is the server name
- L:\mountdir is the component
- FULL is the backup type
- LOCAL is the backup destination
- MC1Q11 is the management class

Binding backups to a policy

You can add, update, delete, or change the processing order of binding statements.

About this task

A backup policy determines how backups on local shadow volumes are managed and retained.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the navigation tree, click **IBM Spectrum Protect**.
3. Select an **Exchange Server**, **SQL Server**, or **File System** instance.
4. In the **Actions** pane, click **Properties**.
5. From the list of available property pages, select **VSS Policy Binding**.
6. Add, update, delete, or change the processing order of existing binding statements.

Tip: You can use an asterisk (*) as a wildcard character to represent all characters.

For example, in the **Server** field, enter the asterisk to bind the policy to all Exchange Servers, all SQL Servers, or all custom application and file system data.

7. To change the processing order, use **Move Up** and **Move Down**.

Policies are processed from the bottom to the top of the file, and processing stops at the first match.

Tip: To ensure that more specific statements are processed before general statements, list the more general specification before the more specific statement.

8. Save the binding statement.

9. Verify new or updated policies and bindings.

a) Run one or more test backup operations.

b) On the **Recover** tab, verify the management classes that are bound to the test backups.

VSSPOLICY statements for backup types

For VSS backups, VSSPOLICY statements are used to associate VSS backups with management classes. When you change from legacy backups to VSS backups, consider the VSSPOLICY statements that you set for the backup.

The VSSPOLICY statements are in a configuration file, for example, `tdpexc.cfg` and `fcmcfg.xml`. A configuration file can include multiple VSSPOLICY statements. The configuration file is read from the bottom to the top of the file. VSSPOLICY statements in the `tdpexc.cfg` file are similar to the INCLUDE statements that are specified in the IBM Spectrum Protect backup-archive client in the `dsm.opt` file.

If no VSSPOLICY statements are included in the configuration file, or if the VSSPOLICY statements do not match the type of backup that is created, the default management class for the policy domain is used. Backup expiration parameters for the default management class might differ from the settings that are used for preexisting legacy backups. For example, the backup expiration period might be set to 30 days. This setting means that after 30 days, the backup is deleted. Verify that the backups expire according to the business needs of your environment.

If you change the `tdpexc.cfg` or `fcmcfg.xml` file, you must restart the IBM Spectrum Protect client acceptor daemon (CAD), IBM Spectrum Protect remote client agent (DSMAgent), and the IBM Spectrum Protect Scheduler Service for Exchange Server. If the DSMAgent service state is set to **Manual (Started)**, stop the service. The DSMAgent service starts when a VSS backup is initiated, but if the service is started and you change the policy settings, the policy settings do not take effect until you restart the service.

Sample VSSPOLICY statements

The following example shows the syntax of a VSSPOLICY statement:

```
VSSPOLICY srv_name db-name backup-type backup-dest mgmtclass
```

where:

- *srv_name* defines the Exchange Server name. You can enter an asterisk (*) as a wildcard character to match all Microsoft Exchange Servers.
- *db-name* defines the database name. You can enter an asterisk (*) as a wildcard character to match all Microsoft Exchange Server groups. Because the name can include a space, use the quotation marks to encapsulate the database name.
- *backup-type* defines the backup type for example, FULL or COPY. You can enter an asterisk (*) as a wildcard character to match all backup types.
- *backup-dest* defines the backup destination. Use the TSM option to back up data to IBM Spectrum Protect, the LOCAL option to back up data to a local disk, enter an asterisk (*) as a wildcard character to match both backup destinations.
- *mgmtclass* defines the IBM Spectrum Protect management class that is used to bind the types of specified backups.

In the following example, the VSSPOLICY statement is commented out. Before you can use this policy statement, you must remove the asterisk character (*) from the first column of each line.

```

-----
* Sample VSSPOLICY Statements
* -----
* These statements are used to bind VSS backup to specific IBM
* Spectrum Protect Server management classes. Adjust the statements
* to meet your needs and remove the leading asterisks to make them
* operational.
* Note: Matching of these policy bindings are from the bottom up.
*****
* Server      Database      Name      BU Type   BU Dest.   Mgmt Class
* -----
VSSPOLICY *          *          FULL      TSM        IUG_TSM
VSSPOLICY *          *          COPY      TSM        IUG_TSM_COPY
VSSPOLICY *          *          COPY      LOCAL      IUG_COPY
VSSPOLICY *          *          FULL      LOCAL      IUG_LOCAL
VSSPOLICY *          "HR"      FULL      LOCAL      MCLASS3
VSSPOLICY SERVER1 "ACT" *      LOCAL      MCLASS2
VSSPOLICY SERVER1 "SG 1" *      TSM        IUG1
*****

```

In the preceding example, the following policy rules are specified:

- Any VSS backups of the *SG 1* database on the Exchange Server *SERVER1* to IBM Spectrum Protect are bound to the management class *IUG1*.
- Any VSS backups of the *ACT* database on the Exchange Server *SERVER1* to *LOCAL* are bound to the management class *MCLASS2*.
- Full VSS backups of the *HR* database on any Exchange Server to *LOCAL* are bound to the management class *MCLASS3*.
- Full VSS backups of any other database on any other Exchange Server to *LOCAL* are bound to the management class *IUG_LOCAL*.
- Copy VSS backups of any other database on any other Exchange Server to *LOCAL* are bound to the management class *IUG_COPY*.
- Copy VSS backups of any other database on any other Exchange Server to IBM Spectrum Protect are bound to the management class *IUG_TSM_COPY*.
- Full VSS backups of any other database on any other Exchange Server to IBM Spectrum Protect are bound to the management class *IUG_TSM*.
- Any type of backup matches a rule because of the wildcard VSSPOLICY statements at the top of the file. Use these types of statements so that you explicitly state what management class is used.

Managing Exchange Database Availability Group members by using a single policy

Online copies of Microsoft Exchange Server databases are maintained in a Database Availability Group (DAG) environment for high availability. To reduce the number of backups on an IBM Spectrum Protect server, set up IBM Spectrum Protect Snapshot to back up database copies from different DAG members under a single DAG node.

About this task

You can prevent IBM Spectrum Protect Snapshot from backing up each database copy separately by backing up the database copies under a single DAG node.

All database copies can be managed as a single entity regardless of where the database copies are backed up from, and whether the backup copies are active or passive at the time of the backup. You can set up a minimum interval between database backups, which ensures that the database copies are not backed up at the same time or backed up too frequently.

Procedure

1. Use the IBM Spectrum Protect configuration wizard to configure the DAG node. Ensure that all the DAG members are configured with the same DAG node name.
 - For VSS backups to IBM Spectrum Protect, specify a node name in the **DAG Node** field on the **TSM Node Names** page in the wizard. This node is used to back up all the DAG.

- For a stand-alone configuration, complete the following steps:
 - In Microsoft Management Console (MMC), select **Exchange workload** and click **Properties**.
 - Click **General**, and specify a node name in the **Back up DAG databases to common node** field.
- 2. Grant permission to the DAG member server to act as a proxy for the DAG node. Issue the **proxynode** command for each member server in the DAG.
For example, issue the following commands:

```
register node backup_archive_client_node password
userID=backup_archive_client_node
```

```
register node data_protection_node password userID=data_protection_node
```

```
grant proxynode target=data_protection_node agent=backup_archive_client_node
```

```
register node DAG_node password userID=DAG_node
```

```
grant proxynode target=DAG_node agent=backup_archive_client_node
```

```
grant proxynode target=DAG_node agent=data_protection_node
```

Tip: If you do not use the configuration wizard to configure the IBM Spectrum Protect server, you must define the proxies and assign proxynode authority to the backup-archive client node and the Data Protection node.

3. For a stand-alone configuration, ensure that the DAG node and the IBM Spectrum Protect Snapshot node are in the same policy domain.

Select a workload from the **Protect and Recover Data** node in the navigation tree of Microsoft Management Console (MMC), and click **Properties** in the **Actions** pane. Set and verify the policy domain name in the **VSS Policy Binding** property page for the Exchange Server workload.
4. Create a backup schedule and specify the **/MINIMUMBACKUPINTERVAL** parameter with the **backup** command.

For example, to back up one copy of a database that contains multiple copies, complete the following steps:

- a) Create a command script that is named C:\BACKUP.CMD by issuing this command:

```
TDPEXCC BACKUP DB1 FULL /MINIMUMBACKUPINTERVAL=60
```

- b) Copy the BACKUP.CMD file to all the DAG members.
- c) Create one schedule and associate all the nodes with this schedule.

5. Run the schedule by using the IBM Spectrum Protect scheduler.

When the backup schedule runs, the minimum backup interval is observed and only one backup is created.

What to do next

To decrease the load on the production Exchange Server, you can specify that the backups are taken from a valid passive database copy. If a valid passive copy is not available, the backup copy is created from the active copy of the database. To add this specification, specify the **/PREFERDAGPASSIVE** on the **backup** command, for example:

```
TDPEXCC BACKUP DB1 FULL /MINIMUMBACKUPINTERVAL=60 /PREFERDAGPASSIVE
```


Data restore overview

IBM Spectrum Protect Snapshot can use the Microsoft Volume Shadow Copy Service (VSS) framework to complete fast and instant restores of database backups. You also restore VSS backups to an alternate database and complete Exchange mailbox restore operations. For SQL Server, you can run legacy restore operations from the IBM Spectrum Protect server.

In a VSS restore operation, you can restore one or more databases from a VSS backup on IBM Spectrum Protect server storage to the original location on the Exchange or SQL Server.

For SQL Server restore operations, you can also restore from a VSS backup to an alternate SQL Server instance. This instance can reside either on the same server where the snapshot is taken or on a different server. For local backups, you can restore only to alternate SQL Server instances on the same server.

VSS fast restore processing

A VSS fast restore operation restores data from a local snapshot. A VSS fast restore operation overwrites any files that exist at the time of the snapshot on the original source location. The file is overwritten with the version that is stored on the snapshot. Data is overwritten even if a file is marked read-only.

You can use VSS fast restore operations for the following tasks:

- Restore Exchange Server VSS backups, full, copy, incremental, and differential backup types.
- Restore data at the database level. However, for custom application and file system data, you can restore data at the file system level only. File overwriting occurs even if the file is marked read-only.
- Restore SQL Server VSS backups to the same SQL Server instance.
- Restore SQL Server VSS backups to an alternate location.
 - You can restore to a database with the same name on either an alternate SQL Server instance on the same server or to an instance on a different server by using the **/fromsqlserver** option. For local backups, you can restore only to alternate SQL Server instances on the same server.
 - You can also restore to an alternate SQL Server instance on the same server or to an instance on a different server where the database does not exist by using the **/relocateddir** option.
- Restore one or more databases or file systems from a VSS snapshot backup on local shadow volumes that are managed by IBM Spectrum Protect Snapshot.
- For SQL, custom application, and file system data, restore the data in Microsoft Windows failover clustering environments.
- In an IBM Spectrum Protect configuration, restore local database backups to only the system that created the backup.

VSS instant restore processing

In an *instant restore* operation, a volume-level snapshot of a local Volume Shadow Copy Service (VSS) backup is restored. The VSS backup must exist on SAN-attached volumes.

A Microsoft VSS instant restore operation restores data by using a hardware-assisted restore method.

Data must be restored on a storage system that runs VSS software and can process VSS instant restore operations. Instant restore processing requires a VSS hardware provider to restore data at the volume level.

Typically, you can restore local VSS backups of SAN-attached volumes from the following storage systems:

- IBM System Storage DS8000 series
- IBM System Storage SAN Volume Controller
- IBM Storwize
- IBM XIV Storage Systems

In Microsoft Windows Server 2012 or later, you can run instant restore operations on IBM and non-IBM storage devices that are enabled to work with a transportable snapshot. A *transportable snapshot* is a cloned image of an Exchange or SQL server database backup that is stored in a SAN cloud, and can be moved from one system to another.

For instant restore operations, use the devices that are listed here: [Storage Architecture Support for Tivoli® Storage FlashCopy Manager and IBM Spectrum Protect Snapshot \(http://www.ibm.com/support/docview.wss?uid=swg21455924\)](http://www.ibm.com/support/docview.wss?uid=swg21455924).

If data is not on an XIV, SAN Volume Controller, or Storwize family systems with space-efficient target volumes, you must ensure that background copies that use the volumes are restored.

You can manually disable VSS instant restore processing so that IBM Spectrum Protect Snapshot uses VSS fast restore processing. Instant restore processing is automatically disabled for these VSS restore operations:

- Restore data to an alternate location
- Restore data to an Exchange Server recovery database
- Restore files by issuing the **restorefiles** command

Restriction: With VSS instant restore processing, SQL Server VSS backups can be restored only to the same SQL Server instance from which they are backed up.

Even though Exchange Server data is restored relatively quickly, the transaction logs must be replayed after a restore operation. The time of recovery for the Exchange Server database increases as the number of logs to be replayed increases.

VSS backups that are restored to alternate databases

IBM Spectrum Protect Snapshot can restore an Exchange Server database backup or DAG active or passive database copy backup, to a recovery database or to an alternate (or relocated) database. For SQL Server restore operations, you can also restore from a VSS backup to an alternate SQL Server instance. This instance can reside either on the same server where the snapshot is taken or on a different server.

Exchange Server VSS backups

This restore method is called *restore into*. If you are restoring a relocated database, use the *restore into* function. You must specify the same database name as the one you are restoring.

Note: If you use the *restore into* function, VSS instant restore capability is automatically disabled.

Backups to local shadow volumes can be restored only to the system where the backups are created.

SQL Server VSS backups

You can restore SQL Server VSS backups either to the same SQL Server instance where the backup was taken or to an alternate SQL Server instance. The following restore options are available:

- Restore local and IBM Spectrum Protect VSS backups. You can restore either to an alternate SQL server instance on the same server where the backup is taken or to a different server. For local backups, you can restore only to alternate SQL Server instances on the same server.
- Restore from VSS backups on virtual machines. You can restore to an alternate SQL server instance on the same server or to an instance on a different server.

For more information, see [Restoring an SQL Server database to an alternate instance](#).

Exchange mailbox restore operations

By using IBM Spectrum Protect Snapshot, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

Restriction:

In the **Mailbox Restore** view of Microsoft Management Console (MMC), you can restore mail items to Unicode and non-Unicode .pst files. In the **Mailbox Restore Browser** view, however, you can restore mail items to only non-Unicode .pst files.

The Recoverable Items folder is a storage area in a mailbox that contains operational data about the mailbox. Depending on the Exchange Server functions that you enabled for the mailbox, you can recover and restore the following types of mail items in the Recoverable Items folder:

- Deleted item retention
- Single item recovery
- In-place hold
- Litigation hold
- Mailbox audit-logging
- Calendar logging

Typically, a mailbox is set up to protect mail items from being accidentally or maliciously deleted, or to retrieve mail items during litigation or investigations.

Mailbox-enabled functions

You can verify whether a mailbox is enabled for mailbox restore operations by running the following Exchange Powershell cmdlets. In the examples, the mailbox is for George Clark:

Deleted item retention

```
Get-Mailbox "george clark" | FL RetentionHoldEnabled,  
RetainDeletedItemsFor, RetainDeletedItemsUntilBackup
```

Single item recovery

```
Get-Mailbox "george clark" | FL SingleItemRecoveryEnabled
```

In-place hold

```
Get-Mailbox "george clark" | FL InPlaceHolds
```

Litigation hold

```
Get-Mailbox "george clark" | FL LitigationHoldEnabled
```

Mailbox audit-logging

```
Get-Mailbox "george clark" | FL AuditEnabled,  
AuditLogAgeLimit
```

Calendar logging

```
Get-Mailbox "george clark" | FL CalendarVersionStoreDisabled
```

Mail items in the Recoverable Items folder

In the mailbox restore views in Microsoft Management Console (MMC), you can recover and restore mail items from the subfolders within the Recoverable Items folder. You can also complete this task by issuing the **restoremailbox** command. The following table lists the subfolders that are included in the Recoverable Items folder.

Table 4. Exchange Server Recoverable Items folder contents

Recoverable Items subfolder	Mailbox-enabled functions	Subfolder contents
Deletions	Deleted item retention	Contains mail items that a user deleted from the Deleted Items folder in their mailbox
Versions	<ul style="list-style-type: none"> • In-place hold • Litigation hold • Single item recovery 	Contains the original and modified copies of the deleted mail items
Purges	<ul style="list-style-type: none"> • Litigation hold • Single item recovery 	Contains all mail items that a user <i>hard deleted</i> , that is, purged from their mailbox
Audits	Mailbox audit-logging	Contains audit log entries
Discovery Holds	In-place hold	Contains mail items that are to be protected from deletion and match <i>hold</i> query parameters
Calendar Logging	Calendar logging	Contains calendar changes that occur within a mailbox

Restriction:

- You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the contents of the email folders.
- You cannot create a subfolder in the Recoverable Items folder in a mailbox.
- You can restore the Recoverable Items content for a public folder mailbox but not for each public folder in the public folder mailbox.

Related tasks

Restoring mailbox data

IBM Spectrum Protect Snapshot backs up mailbox data at the database level, and also restores individual mailbox items from the database backup.

Restoring mailbox messages interactively with the Mailbox Restore Browser

You can use the **Mailbox Restore Browser** to interactively restore a mailbox or items from a mailbox on an Exchange Server.

Configuring your system for mailbox restore operations

To use IBM Spectrum Protect Snapshot for Exchange Server to restore mailboxes and mailbox items, you must complete the configuration wizard.

Related reference

Restoremailbox command

To restore mailbox-level data or mailbox-item-level data from IBM Spectrum Protect Snapshot for Exchange Server backups, use the **restoremailbox** command.

Related information

IBM System Storage requirements

If you use IBM System Storage DS8000 series, SAN Volume Controller, or Storwize family storage systems, be aware of database, log, file, and LUN settings.

Follow these guidelines when you plan for IBM System Storage:

- Place database files on a separate and dedicated logical volume.
- Place logs on a separate logical volume.
- Ensure that the database LUNs are dedicated to only one database or application.

IBM Spectrum Protect Snapshot with IBM SAN Volume Controller and IBM Storwize V7000

The way in which you configure the VSS hardware provider and how you configure your IBM SAN Volume Controller and IBM Storwize V7000 device dictates the characteristics of your VSS snapshots. VSS snapshots are implemented as flashcopys on an IBM SAN Volume Controller and IBM Storwize V7000 device. FlashCopys consist of a source volume and a target volume.

VSS snapshot characteristics differ depending on if any of the following are configured.

- Incremental snapshots
- Space efficient FlashCopys
- Background copy rates

The VSS hardware provider that you use with IBM SAN Volume Controller and IBM Storwize V7000 has the following characteristics when used with incremental snapshots.

- Each VSS snapshot request triggers an incremental refresh of the same target volume; only one FlashCopy pair exists. For this reason, you can keep only one backup version.
- When you delete the VSS snapshot, it is removed from the VSS inventory only. The incremental FlashCopy pair will still exist on the storage device. If you create another VSS snapshot of the same source volume, the process results in an incremental refresh of the target volume.

The following guidelines apply when you use IBM Spectrum Protect Snapshot with SAN Volume Controller- based storage:

- Do not use a combination of space-efficient and fully allocated target volumes. Choose to use either space-efficient or fully allocated volumes for snapshot targets.
- Provision enough target volumes in the SAN Volume Controller VSS_FREE volume group for the backup versions you require.
- The capacity size of the volumes in the VSS_FREE volume group must match the size of the source volumes.
- Do not use a combination of persistent and nonpersistent VSS snapshots.
- Do not mix COPY and NOCOPY snapshot relationships from the same source volume or volumes.
- Enable the autoexpand option for the space-efficient target volumes to avoid out-of-space conditions.
- If you are using fully allocated target volumes and keeping more than one backup version, the IBM SAN Volume Controller or IBM Storwize V7000 device creates multiple target FlashCopy mappings. A mapping might stay in the copying state after all of the source data is copied to the target. This situation can occur if mappings that started earlier and use the same source disk are not yet fully copied. To avoid having dependencies between FlashCopy mappings, try to schedule local backups for IBM SAN Volume Controller and IBM Storwize V7000 storage systems at intervals that are greater than the time required for the background copy process to complete.

Note: You can use VSS instant restore operations with IBM Spectrum Protect Snapshot when multiple backup versions exist.

- Restore from space-efficient target volumes and from fully allocated target volumes for which the background copy of the VSS backup is not completed is supported.

Related tasks

[Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000](#)

If you experience VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems, investigate the configuration in your environment.

Thin provisioning support

Thin provisioning is used to define a storage unit (full system, storage pool, volume) with a logical capacity size that is larger than the physical capacity assigned to that storage unit. A thin-provisioned volume is typically considered a space-efficient (SE) volume.

With SE volumes, you can retain multiple FlashCopy images of a source volume as backups at a much reduced storage cost. You do not need to allocate the full size of the source volume for each backup that is generated.

General guidelines when using SE volumes

Allocate enough space for space-efficient target volumes to hold 120 % of the data that is expected to change on the source volume in the time between snapshots. For example, if a database changes at a rate of 20 % per day, VSS backups complete every six hours, and a steady rate of change throughout the day is assumed. The expected change rate between snapshots is 5 % of the source volume (20/4). Therefore, the allocated space for the space-efficient target volumes is to be 1.2 times 5 % equal to 6 % of the source volume size. If the rate of change is not consistent throughout the day, allocate enough space to the target volumes to accommodate the highest expected change rate for the period between snapshots.

Tip: You can transition your data from fully allocated targets to space-efficient targets by using fully allocated targets as if they are space-efficient when the background copy rate is set to zero.

For SE target volumes, the IBM SAN Volume Controller and IBM Storwize V7000 hardware architectures minimize the space that is required to maintain multiple snapshots of the same source volume. Target volumes are placed into a cascade where each target depends on changes that are recorded in target volumes of subsequent snapshots. For example, assume that four VSS snapshots of a source volume are created. S is the source and T1 through T4 are the targets. T1 is the first, chronologically, and T4 is the last. The following cascade occurs: S -> T4 -> T3 -> T2 -> T1.

With this type of cascade relationship, a copy-on-write process is needed only between the source volume and the last FlashCopy target. Any blocks that remain unchanged on the source volume are not copied. To enable cascading FlashCopy mappings, set the IBM VSS provider background copy value to zero by entering the **IBMvcfg set background Copy 0** command.

To activate the changes, restart the IBM VSS system service after you enter the command.



Warning: When you use the target volumes as backup versions that are managed by IBM Spectrum Protect Snapshot, you must carefully consider the cascaded relationship where multiple SE target volumes have the same FlashCopy source. Some operations, for example, restoring or deleting a snapshot, can result in the unexpected removal of targets, and as a result break the cascaded relationship.

- Delete Example: If T3 was deleted from the chain, then the snapshots that are associated to T2 and T1 would be removed.
- Restore Example: If T3 was restored, the snapshots that are associated to T3 and any future would be removed.

When using space-efficient target volumes, adhere to the following guidelines to avoid the examples listed above.

- Do not delete snapshots manually.

Allow IBM Spectrum Protect Snapshot to delete backup versions that are based on the defined policy to ensure that deletion occurs in the correct order.

- If you delete a local snapshot that is stored on a space-efficient volume (SEV) that has multiple dependent targets, delete the snapshots in the same order in which you created them. You must delete the oldest one first, followed by the second oldest.
- If you perform an instant restore, be aware that the restored versions and newer versions will be removed.



Warning: If FlashCopy mappings are removed from the storage device because of breaks in the cascading FlashCopy mappings chain, the MMC GUI will not reflect those removals immediately.

Some operations, for example, restoring or deleting a snapshot, can result in the unexpected removal of targets, and break the cascaded relationship.

- **Delete Example:** If T3 was deleted from the chain, then the snapshots that are associated to T2 and T1 would be removed. (In this case, the backups associated with T2 and T1 would still show in the MMC GUI until they were expired by policy or deleted.)
- **Restore Example:** If T3 were restored, then the snapshots that are associated to T3 and any future would be removed. (In this case, the backups associated with T3 and any future would still show in the MMC GUI until the next query operation.)

Failover clustering and AlwaysOn Availability

In an SQL Server cluster environment, two different kinds of clustering are supported; AlwaysOn Failover Cluster Instances (FCI) and AlwaysOn Availability Groups (AAGs). An SQL AlwaysOn failover cluster instance provides high availability and disaster recovery at the SQL Server level. AlwaysOn Availability Groups (AAG) provide high availability and disaster recovery at SQL database level. IBM Spectrum Protect Snapshot for SQL Server protects availability databases in both AlwaysOn failover cluster instances and in an AAG.

An AlwaysOn node manages backups of availability databases. This node is a shared node that allows data backups and restores of availability databases from any database replica in the cluster. IBM Spectrum Protect Snapshot for SQL Server treats a backup as originating on a single SQL Server regardless of which node of the cluster is backed up.

Failover cluster instances

In a Windows failover cluster instance with multiple SQL Server instances, the storage is shared and can be accessed by all systems in the cluster. However, only one server in the cluster runs SQL Server services at any given time. When you run a backup, the backup runs on the same server in the cluster that is running the SQL Server service. Therefore, when you run a backup, for example using the Windows scheduler, it must run on this same server in the cluster where the SQL Server instance is active.

AlwaysOn availability groups (AAGs)

An AAG can contain a set of primary databases and multiple copies of the set of primary databases, called secondary databases. Databases in an availability group are called availability databases, and they fail over together as a group. Unlike a failover cluster, in an AAG, storage is not shared because SQL Server uses log shipping to replicate data from the primary database to the secondary database instances.

You can use AAGs with SQL AlwaysOn failover cluster instances to complete the following tasks:

- In an AAG, you can deploy a group of single or clustered server instances, each holding a copy of all databases
- You can use synchronous and asynchronous replication
- You can use log shipping. When a transaction occurs on the primary database, it is shipped to the secondary databases.
- You can use automatic and manual failover modes

Cluster setup considerations for AAGs

To set up AAGs in a Microsoft Windows failover cluster environment or in a Veritas cluster server cluster environment, follow these guidelines:

- Install IBM Spectrum Protect Snapshot for SQL Server on each cluster node and configure each node identically. Specify identical configurations in the IBM Spectrum Protect Snapshot for SQL Server options file.

- Ensure that each availability replica of an availability group is on a different node in the same Windows failover cluster environment.
- Use the Configuration Wizard to register an AlwaysOn node on the IBM Spectrum Protect server. To do so manually, issue the **register node** command on the IBM Spectrum Protect server.
- To access a clustered SQL Server, identify the virtual server name and specify that name in IBM Spectrum Protect Snapshot for SQL Server.
- If you use the IBM Spectrum Protect scheduler to automate data backups, install the scheduler service on each node of the cluster to enable failover support.
- You cannot restore a VSS backup to an alternate instance. Restore VSS backups on the same SQL Server instance where the snapshot is taken.

Tip: VSS and legacy full backups of availability databases on secondary replicas are copy-only. The copy-only option is not automatically used with log backups because you can run log backups that truncate logs on secondary replicas.

Availability database backup operations

IBM Spectrum Protect Snapshot backs up each availability database as a single object, regardless of which availability replica is used for backup and restore operations.

An AlwaysOn Availability Group (AAG) requires SQL Server instances on Windows Failover Cluster nodes. An availability group can have a number of replicas. For example, availability group 1 might have replicas node1, node2, and node3.

A cluster node might be a replica for one or more availability groups. For example, node1 might be a replica for availability group 1 and another availability group.

The AlwaysOn Node is used to manage backups of availability databases. When you work in an IBM Spectrum Protect environment, the AlwaysOn Node is to be common in a Windows Failover Cluster. This presence enables the management of backups of an availability database in a single location, regardless of the replica that is used to complete the backup.

You can run the following types of VSS backup operations:

- Full VSS backups of the primary availability replica
- VSS copy-only full backups of availability replicas

Restriction: The following restrictions apply during availability database backup operations:

- Microsoft does not support legacy full backups on secondary replicas. However, IBM Spectrum Protect Snapshot for SQL Server, does permit you to run a `full` backup of a secondary replica based on IBM Spectrum Protect policy.
- If you use Microsoft SQL Server Standard Edition, Microsoft does not support backups of secondary replicas in an AAG. To back up secondary replicas in an AAG, you must use SQL Server Enterprise Edition. For information, see [Basic Availability Groups \(Always On Availability Groups\)](#).
- When you run a `full` legacy backup of a secondary replica, the underlying implementation of IBM Spectrum Protect Snapshot for SQL Server is to back up the data as `copyfull`. However, IBM Spectrum Protect Snapshot for SQL Server detects the intended full backup operation and applies the IBM Spectrum Protect policy that is associated with the `full` backup type.
- Microsoft Management Console (MMC) and CLI views honor the IBM Spectrum Protect policy that applies to the backup type and in this instance, show the backup type as `full`. For information, see [Active Secondaries: Backup on Secondary Replicas \(AlwaysOn Availability Groups\)\(https://msdn.microsoft.com/en-us/library/hh245119.aspx\)](https://msdn.microsoft.com/en-us/library/hh245119.aspx).

For all backup operations of secondary availability replicas, the secondary replicas must be in the synchronized or synchronizing state.

To assist you with scheduling and load balancing, scheduled backup preference settings of availability groups are also available.

Availability database restore operations

Depending on how you back up availability databases, legacy restore and VSS restore operations are available to restore the availability databases on primary or secondary availability replicas.

Certain restrictions apply when you restore availability databases:

Legacy restore

You can restore an availability database on either a primary or secondary replica.

During the restore process, the restored database is removed from the availability group. When a database is removed from the availability group, the database becomes a local database on that replica. The database is restored as a local database. After the database is restored, you must verify that the data on all replicas is transactionally consistent.

To verify that the data is transactionally consistent, verify that the backup copy contains data and transaction log records. Full backups and differential backups contain data and transaction log records so that the restored database is transactionally consistent.

After you verify that the data is transactionally consistent, manually add the database to the availability group.

VSS restore

You can restore SQL Server VSS backups either to the same SQL Server instance where the snapshot is taken or to an alternate SQL Server instance.

AlwaysOn availability databases

For AlwaysOn availability databases, you must set up IBM Spectrum Protect Snapshot to use an AlwaysOn node name. By default, the AlwaysOn node name is set to the cluster node name for the Availability Group in SQL Server 2012, and later versions.

Automated IBM Spectrum Protect server failover for data recovery

If you use IBM Spectrum Protect Snapshot with the IBM Spectrum Protect configuration, IBM Spectrum Protect Snapshot can automatically fail over to a secondary server for data recovery when there is an outage on the IBM Spectrum Protect server.

The IBM Spectrum Protect server that IBM Spectrum Protect Snapshot connects to for backup services is called the *primary server*. If the primary server is set up for node replication, the client node data on the primary server can be replicated to another IBM Spectrum Protect server, which is the *secondary server*.

Depending on your configuration, you must set up the following nodes for replication on the primary server:

- IBM Spectrum Protect Snapshot node
- VSS Requestor node
- IBM Spectrum Protect Remote Client Agent (DSMAGENT) node for offloaded backups to the primary server
- Exchange Server Database Availability Group (DAG) node for backups of databases in a DAG
- AlwaysOn node for backups of availability databases in an AlwaysOn Availability Group (AAG) on SQL Server 2014 and later versions. The AlwaysOn node is a shared nodes that facilitates backups and restores of availability databases from any replica.

During processing, connection information for the secondary server is automatically sent to IBM Spectrum Protect Snapshot from the primary server. The secondary server information is saved to the client options `dsm.opt` file.

Each time the backup-archive client logs on to the server for backup services, it attempts to contact the primary server. If the primary server is unavailable, the backup-archive client automatically fails over to the secondary server. In failover mode, you can restore data that is replicated to the secondary server.

When the primary server is online again, the backup-archive client automatically fails back to the primary server the next time the backup-archive client connects to the server.

Requirements: To ensure that automated client failover can occur, IBM Spectrum Protect Snapshot must meet the following requirements:

- IBM Spectrum Protect Snapshot must be at least at V4.1 level or later.
- The primary server, secondary server, and backup-archive client must be at least at V7.1.1 level or later.
- The primary and secondary servers must be set up for node replication.
- The following nodes must be configured for replication with the `replstate=enabled` option in each node definition on the server:
 - IBM Spectrum Protect Snapshot node
 - VSS Requestor node
 - Remote DSM agent node for offloaded backups
 - DAG node, if applicable
 - AlwaysOn node, if applicable
- Before the connection information for the secondary server can be sent to IBM Spectrum Protect Snapshot, the following processes must occur:
 - You must back up data at least one time to the primary server.
 - The following nodes must be replicated at least one time to the secondary server:
 - IBM Spectrum Protect Snapshot node
 - DAG node, if applicable
 - AlwaysOn node, if applicable

Restriction:

The following restrictions apply to IBM Spectrum Protect Snapshot during failover:

- Any operation that requires data to be stored on the IBM Spectrum Protect server, such as backup operations, are not available. You can use only data recovery functions, such as restore or query operations.
- Schedules are not replicated to the secondary server. Therefore, schedules are not run while the primary server is unavailable.
- If the primary server stops before or during node replication, the most recent backup data is not successfully replicated to the secondary server. The replication status of the file space is not current.



Attention: If you restore data in failover mode and the replication status is not current, the recovered data might be corrupted. You must wait until the primary server comes back online before you can restore the data.

Chapter 2. Planning

You can install and configure IBM Spectrum Protect Snapshot software on a local system or on a virtual machine. From one IBM Spectrum Protect Snapshot installation, you can manage all of the IBM Spectrum Protect Snapshot installations in your enterprise.

About this task

Before you implement your backup and restore strategies, review the security requirements and other guidelines that are specific to protecting data in your IBM Spectrum Protect Snapshot for Windows environment. Consider how to manage your IBM Spectrum Protect policy, and set IBM Spectrum Protect Snapshot configuration options and preferences.

Storage capacity requirements

With IBM Spectrum Protect Snapshot, you need storage space for the product installation, space to store IBM Spectrum Protect Snapshot metadata, and space on the storage device for the snapshot backups.

Product installation

The space that you need for the product installation of IBM Spectrum Protect Snapshot depends on the components that are installed. Space requirements also depend on required maintenance updates and required operating systems, applications, and other software currency support.

When you plan a product installation, the following components is required depending on the data that you want to protect:

- Microsoft Management Console (MMC) and the VSS Requestor are required components. You install MMC when you install the software by running the `setupfcm.exe` file. The VSS Requestor is automatically installed for all installations.

IBM Spectrum Protect Snapshot metadata

IBM Spectrum Protect Snapshot uses disk space to store product metadata that tracks and manages snapshots (point-in-time consistent copies of application data).

The amount of space that is required is directly proportionate to the number of snapshots that you maintain on the system. For each snapshot that you plan to retain on the system, ensure that at least 1 MB of free disk space is available to store the metadata.

The amount of disk space that is required to store metadata depends on the configuration of your environment.

Configuration with only IBM Spectrum Protect Snapshot

If you are protecting Microsoft Exchange Server data, IBM Spectrum Protect Snapshot retains mailbox history information in the metadata to support individual mailbox restore (IMR) processing. The amount of space that is required to store the metadata is proportionate to the number of mailboxes and log files in the entire organization. For each user mailbox in an organization, ensure that at least 50 KB of disk space is available to store the metadata.

Configuration with IBM Spectrum Protect and IBM Spectrum Protect Snapshot

If you are protecting Exchange Server data, IBM Spectrum Protect Snapshot retains the mailbox history information that is stored on the IBM Spectrum Protect server. In this configuration, no disk space is required for IBM Spectrum Protect Snapshot.

Snapshot copies

Snapshot copies of application data require the most storage space. The amount of space that is required depends on the following factors:

- The VSS provider that you use and the configuration of the VSS provider
- The total size of all source volumes that contain the application data
- The rate at which the source volumes are altered after a snapshot is taken

Full snapshots are the standard type of FlashCopy snapshot. On SAN Volume Controller, DS8000, and Storwize V7000 storage devices, full snapshot copies require the same amount of space as the corresponding source volumes. However, with the Windows System VSS provider, space-efficient copies on SAN Volume Controller and the XIV system initially require space for only the metadata. The space requirement for snapshot copies increases with every volume block that changes on the corresponding source volume. As more source volume blocks change, more space is required for the target volumes that represent a snapshot copy of those applications.

Best practices for IBM Spectrum Protect Snapshot with IBM XIV 11.6 Real-time Compression

You can use IBM XIV 11.6 Real-time Compression with IBM Spectrum Protect Snapshot.

About this task

The usage of IBM Spectrum Protect Snapshot with compressed volumes does not change. However, when you transform volumes managed by IBM Spectrum Protect Snapshot from the uncompressed state to the compressed state (or if you transform from compressed to uncompressed), use the following list of behaviors as a guide:

Procedure

1. When source volume transformation is in progress (from uncompressed to compressed, or compressed to uncompressed), most IBM Spectrum Protect Snapshot operations (for example, back up, restore, and mount) fail. The IBM Spectrum Protect Snapshot software returns the **FMV1235E (RC-1)** message.
Perform the volume transformation at a time that does not overlap with scheduled backups or other IBM Spectrum Protect Snapshot actions running on the volume that is being transformed.
2. With the XIV system, you can transform a volume from uncompressed to compressed state (or compressed to uncompressed state) using one of the following options:
 - With the `delete_source=yes` option, delete all volume backups. If you do not delete the volume backups, the transform is unsuccessful. You can use the IBM Spectrum Protect Snapshot GUI or CLI to manually delete the backups before the transform operation runs.
 - With the `delete_source=no` option, the volume backups are retained. After the transform completes, the original (source) volume is hidden from the host system. The original volume is replaced by the transformed volume. Any instant restore operation completed with the backups made before the transformation are restored to the hidden volume on the storage device. The restore is not made to the volume seen by the host. Note that the restore to the volume seen by the host appears to be successful, but the source volume visible to the host system is unchanged.

When using IBM Spectrum Protect Snapshot to protect volumes to be transformed, delete the existing snapshot backups, regardless of the `delete_source` option setting.

Chapter 3. Installing and upgrading

IBM Spectrum Protect Snapshot wizards guide you through the installation, upgrade, and configuration of IBM Spectrum Protect Snapshot. After you complete the setup and configuration wizards, your system is ready to back up and restore data.

IBM Spectrum Protect Snapshot provides the following wizards for installation and configuration tasks:

Standalone configuration wizard

Use this wizard to configure IBM Spectrum Protect Snapshot to manage snapshot backups as a standalone computer.

IBM Spectrum Protect configuration wizard

Use this wizard to configure IBM Spectrum Protect Snapshot in an environment that is integrated with IBM Spectrum Protect. This integration provides data protection and centrally managed, policy-based administration.

Mailbox restore only configuration wizard

Use this wizard to configure IBM Spectrum Protect Snapshot to restore mailboxes from mounted EDB files. Extra data protection features are not available. This configuration is ideal when you have a specific task to complete and do not want the additional IBM Spectrum Protect Snapshot software functions.

Prerequisites

Before you install IBM Spectrum Protect Snapshot, ensure that your system meets the minimum hardware, software, and operating system requirements.

To run data protection operations with IBM Spectrum Protect server, you must install the correct product license in the correct installation directory. If you cannot configure the software, verify that the product license file is correctly installed in one of these directories:

- For IBM Spectrum Protect Snapshot for Microsoft Exchange Server, install the `excclient.lic` license file in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server installation directory.
- For IBM Spectrum Protect Snapshot for Microsoft SQL Server, install the `sqlclient.lic` license file in the IBM Spectrum Protect Snapshot for Microsoft SQL Server installation directory.
- For IBM Spectrum Protect Snapshot, install the `fcclient.lic` license file in the IBM Spectrum Protect Snapshot installation directory.

The installation wizard verifies many of the prerequisites as part of its verification process. However, some prerequisites cannot be automatically verified, for example, the host bus adapter (HBA) or multipath I/O (MPIO) software that is required for your VSS provider.

In addition, the IBM Spectrum Protect Snapshot product comprises multiple components that support different operating systems, databases, and applications. Hardware and software requirements change over time due to maintenance updates and the addition of operating system, application, and other software currency support. Before you begin the installation process, always verify that your environment meets the hardware and software prerequisites.

For more information, review the Hardware and Software Requirements technote that is associated with the level of your IBM Spectrum Protect Snapshot program. This technote is available at this web page: [IBM Spectrum Protect Snapshot - All Requirements Doc \(http://www.ibm.com/support/docview.wss?uid=swg21427692\)](http://www.ibm.com/support/docview.wss?uid=swg21427692). Follow the link to the requirements technote for your specific release or update level.

If you are protecting IBM Spectrum Protect Snapshot databases on Microsoft Windows Server 2008 and later versions, you must install Microsoft Windows PowerShell Version 3.0 or later versions. For more information, see [Microsoft TechNet: Installing Windows PowerShell \(http://technet.microsoft.com/en-us/library/hh847837.aspx\)](http://technet.microsoft.com/en-us/library/hh847837.aspx).

Installation process might require a restart

If you do not install all of the prerequisites before you start the installation process, the installation process might require a restart. As part of the installation process, one or more Microsoft C++ redistributable packages are installed, if they are not already installed on the Windows workstation. These packages can also be automatically updated by the Windows Update service. If the packages are updated, the update can cause the system to restart when you start the installation program.

Additionally, because the Microsoft Visual Studio C++ redistributable package is a shared Windows component, other applications that have dependencies on the package might be stopped or restarted by Windows as part of the installation or upgrade of the C++ redistributable package. Schedule installations and upgrades during a maintenance window when other applications are not be adversely affected if they are stopped or restarted when the C++ redistributable package is installed. Monitor other applications after the installation is complete to see whether any applications were stopped and not restarted.

Virtualization environment resources

For more information about the virtualization environments that you can use with IBM Spectrum Protect Snapshot, see this web page: [IBM Tivoli Storage Manager \(TSM\) and IBM Spectrum Protect™ guest support for Virtual Machines and Virtualization \(http://www.ibm.com/support/docview.wss?uid=swg21239546\)](http://www.ibm.com/support/docview.wss?uid=swg21239546).

Installing IBM Spectrum Protect Snapshot for Windows

The configuration wizard guides you through installing IBM Spectrum Protect Snapshot on your computer.

Before you begin

Before you install and configure IBM Spectrum Protect Snapshot, verify that you satisfy the hardware and software requirements.

IBM Spectrum Protect Snapshot installation packages are delivered electronically through an IBM® download site. You must extract the installation files from the download site.

About this task

IBM Spectrum Protect Snapshot is available in both licensed and maintenance packages. The installation process differs based on the package type.

Licensed package

Includes a license enablement file that is only available from your software distribution channel, such as Passport Advantage®, and includes the initial General Availability release of a product or component.

Maintenance update (fix pack or interim fix package)

Available from the maintenance delivery channel, and can sometimes be used to refresh the software distribution channel. Maintenance packages do not contain license enablement files and must be installed after a licensed package.

See the README . FTP file for information about how to install a fix pack or interim fix package. The README . FTP file is available in the same directory where the maintenance package is downloaded.

Procedure

1. Log on with administrator credentials and complete the following steps:
 - a) Download the appropriate package file from one of the following websites.
 - For a first-time installation or a new release go to Passport Advantage at [IBM Passport Advantage](#). Passport Advantage is the only website from which you can download a licensed package file.

- For a maintenance fix, go to this FTP site and to the directory that contains the maintenance fix version that you require, [Index of Tivoli Storage FlashCopy Manager patch files \(ftp://public.dhe.ibm.com/storage/tivoli-storage-flashcopymanager/patches/\)](ftp://public.dhe.ibm.com/storage/tivoli-storage-flashcopymanager/patches/).

b) After you download the package, complete the following steps:

- Verify that you have enough space to store the installation files when they are extracted from the product package.
- Change to the directory where you placed the executable file.

Tip: In the next step, the files are extracted to the current directory. The path must contain no more than 128 characters. Extract the installation files to an empty directory. Do not extract the files to a directory that contains previously extracted files, or any other files.

- Either double-click the executable file, or enter the following command on the command line to extract the installation files. The files are extracted to the current directory.

```
package_name.exe
```

where package_name is like:

```
version-TIV-FCM-Win.exe
```

For example, for version 8.1.4, the package_name is

```
8.1.4.0-TIV-FCM-Win.exe
```

- Follow the installation instructions that are displayed.
 - Click **Finish** to complete the installation of IBM Spectrum Protect Snapshot.
- Configure IBM Spectrum Protect Snapshot by using the configuration wizard.
 - Start Microsoft Management Console (MMC). Click **Start > All Programs > IBM Spectrum Protect Snapshot > IBM Spectrum Protect Snapshot Management Console**.
When you start MMC, a welcome page is displayed. If IBM Spectrum Protect Snapshot is not configured, the configuration wizard starts.
 - If the configuration wizard does not start automatically, go to **IBM Spectrum Protect > Dashboard > Manage > Configuration > Wizards** in the tree view, and select one of the following wizards:
 - **Standalone configuration wizard**
 - IBM Spectrum Protect **configuration wizard**
 - **Mailbox Restore Only configuration wizard**
 Click **Start** in the **Actions** pane.
 - In the configuration wizard, select the applications to protect, verify requirements, provision, and configure.
 - After you complete the configuration wizard, verify your configuration by selecting each workload instance in the tree view and completing the following steps:
 - Click the **Automate** tab.
 - Click **Open** in the toolbar.
 - Type `verify`. The following three file names are listed.

```
verify_sql.txt
verify_exc.txt
verify_fs.txt
```

The `verify_fs.txt` file is used with MMC and is part of the base product installation.

- Select and open the file that matches the workload.
- Click **Run** in the toolbar.

If the commands run on the command-line interface with no warnings or errors, the configuration is verified.

The `verify_sql.txt` file contains the following commands:

```
query tdp
query tsm
query sql
```

The `verify_exc.txt` file contains the following commands:

```
query tdp
query tsm
query exchange
```

The `verify_fs.txt` file contains the following commands:

```
query component
query config
```

4. Back up and restore a set of test data. Refine your configuration settings as necessary.
5. Define the policy settings and scheduled operations.

Related concepts

[Security requirements for backup and restore operations](#)

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Silently installing IBM Spectrum Protect Snapshot for Windows

You can use the setup program to implement a silent (unattended) installation of IBM Spectrum Protect Snapshot for Windows.

Before you begin

Before you install and configure IBM Spectrum Protect Snapshot for Windows, verify that you satisfy the hardware and software requirements. IBM Spectrum Protect Snapshot installation packages are delivered electronically through an IBM® download site.

Tip: For a first-time installation or a new release, go to Passport Advantage at [IBM Passport Advantage](#). Passport Advantage is the only website from which you can download a licensed package file.

You must install two components: Microsoft Management Console (MMC) and IBM Spectrum Protect client. The setup programs for these components are provided in the installation package file.

IBM Spectrum Protect Snapshot Management Console setup program

(64-bit) `.\fcm\x64\mmc\\enu\spinstall.exe`

About this task

To ensure a consistent configuration and to avoid having 25 different people enter IBM Spectrum Protect Snapshot parameters, an administrator can choose to produce an unattended installation package and make it available to the 25 sites. The installation package can be placed in a shared directory on a file server for distribution across the different sites.

Procedure

1. Enter the following commands to silently install both components to the default installation. The setup program is on the directory where you extracted your installation files.

```
.\fcm\x64\mmc\\enu\spinstall.exe /s /v/qn
```


where *version* is the version of IBM Spectrum Protect Snapshot you want to install.

2. Run the `spinstall.exe` file with the following options. Specify each command on a single line from a **Run as Administrator** command line.

The following examples are commands that specify the target directory, the features, start suppression, and logging.

```
.\fcm\x64\mmc\<<version>\enu\spinstall.exe /s /v"INSTALLDIR=
\"C:\Program Files\Tivoli\"ADDLOCAL=\"Client\" TRANSFORM=1033.mst
REBOOT=ReallySuppress/qn /l*v\"C:\Program Files\Tivoli\FlashCopyManager\logs\
fcm.log\""
```

3. Review these guidelines as you complete the installation process:

- You must place a backslash (\) before each quotation mark that is within an outer set of quotation marks (").
- For a single-line command, press **Enter** only when all the parameters are entered.
- You must place quotation marks (") around the following text:
 - A directory path that contains spaces.
 - An argument that specifies multiple features. Although you must use quotation marks around the complete argument, you must still place a backslash before each internal quotation mark.
- All features that are listed in a custom installation must be listed after the **addlocal** option.
- Setting the **rebootyesno** option to *No* applies only to the installation of the IBM Spectrum Protect Snapshot software. The installation package includes a number of prerequisites that are installed by IBM Spectrum Protect Snapshot. Ensure that all the prerequisites are installed before you start the silent installation, and then set the **rebootyesno** option to *No* to avoid a restart after the silent installation process finishes.

Options in silent installations

The following options can be applied to both silent installation methods, the setup program and the Microsoft Installer (MSI) program.

Review the list of silent installation options in the following tables:

Table 5. Silent installation options

Option	Description
/i	Specifies the program is to install the product.
/l*v	Specifies verbose logging.
/qn	Runs the installation without running the external user interface sequence.
/s	Specifies silent mode.

Table 5. Silent installation options (continued)

Option	Description
/v	Specifies the Setup Program to pass the parameter string to the call it makes to the MSI executable program (msiexec . exe). Note the following syntax requirements when you use the /v option: <ul style="list-style-type: none"> • A backslash (\) must be placed in front of any quotation marks (" ") that are within existing quotation marks. • Do not include a space between the /v command-line option and its arguments. • Multiple parameters that are entered with the /v command-line option must be separated with a space. • You can create a log file by specifying the directory and file name at the end of the command. The directory must exist when you run a silent installation.
/x	Specifies the program to uninstall the product.
addlocal	Specifies features to install.
allusers	Specifies which users can access the installation package.
installdir	Specifies the directory where IBM Spectrum Protect Snapshot for Windows is to be installed.
reboot	Specifies whether to prompt the user to restart the system after silent installation. <p>Force Always prompts user to restart after silent installation.</p> <p>Suppress Suppresses prompt to restart after silent installation.</p> <p>ReallySuppress Suppresses all restarts and prompts to restart after silent installation.</p>
rebootyesno	Specifies whether to restart the system after silent installation. Specify <i>Yes</i> to restart the system after silent installation. Specify <i>No</i> not to restart the system after silent installation.

Setting the **rebootyesno** option to *No* applies only to the installation of the IBM Spectrum Protect Snapshot for Windows software. The installation package includes a number of prerequisites for IBM Spectrum Protect Snapshot for Windows to install if those prerequisite components are not already installed on the system. Ensure that all the prerequisites are installed before you start the silent installation, then set the **rebootyesno** option to *No* so that no system restart is required after the silent installation process finishes.

The following tables list the silent installation features (case-sensitive) that apply to the base client only.

Table 6. Silent installation features (base client only)

Feature	Description
Client	IBM Spectrum Protect Snapshot for Windows code

Table 7. Silent installation features (IBM Spectrum Protect Snapshot for Windows)

Feature	Description
Plug-in	IBM Spectrum Protect (enables LOCAL VSS operations, and offloaded VSS backups)

Creating and testing a silent installation package on a DVD or a file server

The administrator can choose to make an installation package available by burning a DVD or placing the package in a shared directory on a file server.

Before you begin

Before you begin, you must choose a location for the package. If you are burning a DVD, it is convenient to use a staging directory. If you are placing the package on a file server, you can use a staging directory or build the package directly on the file server.

About this task

Typically, the installation package contains the IBM Spectrum Protect Snapshot for Windows code distribution files and a batch file for a silent installation.

Procedure

1. Issue the following commands to create the package:

Table 8. Commands for creating a silent installation package

Command	Description
<code>mkdir c:\tdpdpkg</code>	Create a staging directory for the silent-install package
<code>cd /d c:\tdpdpkg</code>	Go to the staging directory
<code>xcopy g:*.* . /s</code>	Copy the DVD distribution files to the staging directory
<code>copy c:\spinstall.bat</code>	Replace the existing <code>spinstall.bat</code> with the one created in the previous step

This example uses `c:\tdpdpkg` as a staging directory. For more information on creating a silent installation batch file, and for a sample `spinstall.bat` script, see [“Batch files usage in silent installations”](#) on page 36.

2. After you create the installation package, test the silent installation.
3. After you complete the test, place the package on a DVD or make it available from a shared directory.
4. After you make the package available on a DVD or from a shared directory, complete these steps to run the silent installation package on another computer.

Option	Description
From a silent installation package on DVD:	Enable the <code>autostart</code> option to cause the silent installation to begin as soon as the DVD is inserted into the drive. If you do not enable the <code>autostart</code> option, start the <code>spinstall.bat</code> file from the root of the DVD by issuing the following command: <pre>cd /d g:\ spinstall.bat</pre>
From a distribution directory:	If the package is placed in a shared directory that is called <code>tdpdpkg</code> at <code>\machine1\d\$</code> , another computer can run the <code>net use x: \</code>

Option	Description
	<p>\machine1\d\$ command to share the drive as drive x. You can issue the following command:</p> <pre data-bbox="610 268 818 323">cd /d x:\tdpdpkg spinstall.bat</pre>

In either case, the silent installation begins. Allow enough time for the unattended installation to complete. No visual cues exist to inform you when the installation is finished, although you can add visual cues to the batch file.

Batch files usage in silent installations

You can create a batch file to begin the silent installation with the parameters that you want to use.

The following script is a sample script (c:\spinstall.bat) of an unattended installation:

```
@echo off
rem =====
rem sample silent install script
rem
call x:\fcm\x64\mmc\8170\enu\spinstall.exe /s
/v"INSTALLDIR="C:\Program Files\Tivoli\"
ADDLOCAL="Client" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /l*v "C:\Temp\FCM\mmcsinstallLog.txt\"
rem
call x:\fcm\x64\exc\8170\enu\spinstall.exe /s
/v"INSTALLDIR="C:\Program Files\Tivoli\tsm\"
ADDLOCAL="Client" TRANSFORM=1033.mst
REBOOT=ReallySuppress /qn /l*v "C:\Temp\FCM\wspinstallLog.txt\"
rem =====
rem code could be added after the
rem installation completes to
rem customize the dsm.opt files
rem if desired
rem =====
```

Silent installation error messages

The **spinstall.exe** program can produce error messages if it cannot start properly.

In most cases, administrators encounter these messages when a severe error occurs. Users rarely see these messages. When you get an error message, it displays in a message box. Every error message has a number. These messages are system error messages and there is no way to suppress them in your script.

Silently installing IBM Spectrum Protect Snapshot on Windows Server Core

You can use silent installation methods to install IBM Spectrum Protect Snapshot in a Windows Server Core environment, without any user interaction.

About this task

You can use either the setup program or the Windows installer (MSI) program for the unattended installation of IBM Spectrum Protect Snapshot for Windows.

Silently installing IBM Spectrum Protect Snapshot on Windows Server Core with the setup program

You can use the setup program to silently install IBM Spectrum Protect Snapshot on Windows Server Core.

Procedure

For more information, see [“Silently installing IBM Spectrum Protect Snapshot for Windows”](#) on page 32.

Silently installing IBM Spectrum Protect Snapshot on Windows Server Core with the Microsoft Installer program

You can use the Microsoft Installer (MSI) program, `msiexec.exe`, to implement a silent installation of IBM Spectrum Protect Snapshot for Windows. You can also use the MSI program to silently install IBM Spectrum Protect Snapshot on Windows Server Core.

Before you begin

IBM Spectrum Protect Snapshot for Windows must be installed from an account that is a member of the local Administrators group for the system on which IBM Spectrum Protect Snapshot for Windows is running.

Important: Unlike the `spinstall.exe` program, the `msiexec.exe` program does not install any prerequisites. When you use `msiexec.exe`, you must install all prerequisites manually.

Procedure

- To install Microsoft Management Console (MMC), enter each of these **msiexec** commands on a single line from a **Run as Administrator** command line.

```
msiexec /i "x:\fcm\aaa\mmc\<version>\enu\IBM Spectrum Protect Snapshot.msi"  
RebootYesNo="No" Reboot="ReallySuppress" ALLUSERS=1  
INSTALLDIR="c:\program files\tivoli" ADDLOCAL="Client"  
TRANSFORMS="x:\fcm\aaa\mmc\<version>\enu\1033.mst" /qn /l*v "c:\temp\log_fcm.log"
```

where *x*: is your DVD drive and *aaa* is either x86 or x64.

What to do next

Important:

- You must place quotation marks (") around the following items:
 - A directory path that contains spaces.
 - An argument that specifies multiple features. Although you must use quotation marks around the complete argument, you must still place a backslash before each internal quotation mark.
- All features that are listed in a custom installation must be specified after the **addlocal** option.

Upgrading IBM Spectrum Protect Snapshot

You can upgrade IBM Spectrum Protect Snapshot with the latest versions of IBM Spectrum Protect Snapshot for Exchange Server and IBM Spectrum Protect Snapshot for SQL Server.

Before you begin

- Install IBM Spectrum Protect Snapshot. When you extract and install the IBM Spectrum Protect Snapshot `setupFCM.exe` package, ensure that you leave all the source installation binary files on your local system.
- Run the configuration wizard and verify your IBM Spectrum Protect Snapshot version. The configuration wizard does not run if it cannot locate the installation package binary files on your system.

Procedure

- Download the latest patch files for IBM Spectrum Protect Snapshot for Exchange Server or IBM Spectrum Protect Snapshot for SQL Server at [Index of Tivoli Data Protection patches](#).
- Extract the patch files that you downloaded, and run `setupFCM.exe`.

Ensure that you leave all the source installation binary files on your local system or the configuration wizard might not run.

3. To start Microsoft Management Console (MMC), click **Start > All Programs > IBM Spectrum Protect Snapshot > IBM Spectrum Protect Snapshot Management Console**.

The system detects the patch files that you installed for IBM Spectrum Protect Snapshot for Exchange Server or IBM Spectrum Protect Snapshot for SQL Server, and identifies the version.

4. In the Welcome page, click **OK**.

The configuration wizard automatically starts and can vary depending on the software licenses that are found on the system. If the configuration wizard does not start automatically, click **IBM Spectrum Protect Snapshot** in the navigation tree, and click **Configuration**. Then, double-click **Wizards**.

5. In the configuration wizard, select to configure either Exchange Server or SQL Server installed components.

The configuration wizard guides you through the process of provisioning and installing the remaining files for the selected Data Protection component. When the configuration wizard is complete, the Data Protection component version is displayed.

6. At any stage, rerun the configuration wizard to verify the Data Protection component version that IBM Spectrum Protect Snapshot is running.

What to do next

After you upgrade IBM Spectrum Protect Snapshot, you can restore, mount, and unmount any local backups that are created with an earlier version of the software. Use the upgraded version of the software to complete this task. If you use an older version of the software, errors occur.

IBM Spectrum Protect Snapshot migration

You can migrate data from earlier versions of IBM Spectrum Protect Snapshot.

After you upgrade to a newer version of IBM Spectrum Protect Snapshot, use VSS restore for local VSS backups that were originally created with the older version of the software.

If you used a previous version of IBM Spectrum Protect Snapshot in a Microsoft clustering environment, and you upgrade to a newer version of IBM Spectrum Protect Snapshot, any existing backups that are completed on cluster disks do not count toward the maximum number of versions. New backups for clustered disks that are completed with the newer version of IBM Spectrum Protect Snapshot are managed logically for the cluster. Except for the active backup, older backups eventually expire. When you no longer must retain the active backup, the active backup must be deleted by using the **delete backup** command. You can restore the existing backups.

Managing migrated backups to a Database Availability Group node

When you configure IBM Spectrum Protect Snapshot to back up databases in a DAG to a common DAG node, all DAG databases are backed up with the new DAG node name.

Before you begin

If you are migrating from an IBM Spectrum Protect Snapshot version that is earlier than V3.2, manage the backups from the previous versions by following these guidelines:

- Do not mix backups that are created with previous versions of IBM Spectrum Protect Snapshot with new backups that are created by using the DAG node. To separate the backups, keep the previous backups under the previous Data Protection node name that is defined in the `dsm.opt` file in the `C:\Program Files\Tivoli\tsm\TDPEXchange` directory, and use a new DAG node name to store the new backups.
- To view or restore a backup that is stored under the previous node name, you must change the IBM Spectrum Protect Snapshot configuration.
- You must manually delete backups over time if the old backups are no longer useful.

Procedure

1. After you complete your migration, ensure that the first backup you do is a full backup.
2. To view and restore backups that are stored under the previous Data Protection node name, complete these steps:
 - a) Remove the **DAG Node** by using the **General** properties page, configuration wizard, or the **set** command on the command line.
 - b) Restart or refresh Microsoft Management Console (MMC) or command-line interface.
 - c) Click the **Recover** tab in MMC, or run a `tdpexcc query tsm *` command.
Because the **DAG Node** parameter is not set, IBM Spectrum Protect Snapshot lists the backups that are stored under the IBM Spectrum Protect Snapshot node.
 - d) Proceed to restore one or more of the listed backups.
3. Delete the backups that are expired.

Uninstalling IBM Spectrum Protect Snapshot

When you install IBM Spectrum Protect Snapshot, some components are saved to your system. You can remove components by using the Windows **Add or Remove Programs** or **Programs and Features** in the Windows control panel. You must manually remove any remaining files, registry keys, or Windows services that are created by IBM Spectrum Protect Snapshot.

Before you begin

Log into a Windows account with administrator privileges.

About this task

This procedure assumes that a default IBM Spectrum Protect Snapshot configuration is in place.

Use this procedure to completely remove all IBM Spectrum Protect Snapshot data from a computer. Adjust the path in the example to suit your environment.

Procedure

1. Copy any files that you want to keep from the `c:\Program Files\Tivoli` directory and its subdirectories to a different directory.
For example, you might have configuration files that you want to save.
2. Delete any IBM Spectrum Protect Snapshot scheduled tasks:
 - a) Select the **Scheduling** node in the IBM Spectrum Protect Snapshot tree view.
 - b) Select each scheduled task that is listed in the Schedules section of the results pane, and click **Delete**.
3. Stop any IBM Spectrum Protect Snapshot components that are running.
4. Delete any existing IBM Spectrum Protect Snapshot snapshots by issuing the **DELETE BACKUP** command.
5. Enter the following commands. You can use the command `dsmcutil list` to display any IBM Spectrum Protect Snapshot services that are installed.
 - a) `cd /d "c:\program files\tivoli\tsm\baclient"`
If necessary, replace `c:\program files\tivoli` with the correct installation folder.
 - b) `dsmcutil remove /name:"TSM Remote Client Agent"`
Important: Remove the IBM Spectrum Protect Remote Client Agent before you remove the IBM Spectrum Protect Client Acceptor, or the IBM Spectrum Protect Client Acceptor cannot be removed.
 - c) `dsmcutil remove /name:"TSM Client Acceptor"`

6. From the **Control Panel** window, open **Add or Remove Programs** or **Programs and Features**.
7. Uninstall the following items if listed:
 - IBM Spectrum Protect Snapshot
 - IBM Spectrum Protect client
 - IBM Spectrum Protect for Databases: Data Protection for Microsoft SQL Server
 - IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server
8. Find the IBM Spectrum Protect staging files and remove them from the file system. Run the following commands:
 - `cd /d "c:\program files\tivoli"`
If the IBM Spectrum Protect staging files are not in the default locations, manually remove the files. If necessary, replace `c:\program files\tivoli` with the correct installation folder.
 - `rd /s flashcopymanager`
 - `rd /s tsm`
9. Enter the following command:

```
reg query hklm\software\ibm
```

A list of registry keys are displayed. For example:

```
HKEY_LOCAL_MACHINE\software\ibm\ADSM
HKEY_LOCAL_MACHINE\software\ibm\FlashCopyManager
HKEY_LOCAL_MACHINE\software\ibm\GSK7
HKEY_LOCAL_MACHINE\software\ibm\GSK8
```

10. Enter the following commands from a **Run as Administrator** command prompt window.
 - a) Enter this command if you want to completely remove the IBM Spectrum Protect backup-archive client from the system: `reg delete HKLM\SOFTWARE\IBM\ADSM`
You can uninstall IBM Spectrum Protect Snapshot, but continue to use IBM Spectrum Protect backup-archive client.
 - b) `reg delete HKLM\SOFTWARE\IBM\FLASHCOPYMANAGER`
11. Before you enter the following commands, verify these requirements:
 - The entries `HKEY_LOCAL_MACHINE\software\ibm\GSK7` and `HKEY_LOCAL_MACHINE\software\ibm\GSK8` were included in the command output that is shown in Step 9.
 - No other applications are using IBM GSKIT.

If either of the preceding conditions exist, enter the following commands:

 - a) `reg delete HKLM\software\ibm\GSK7`
 - b) `reg delete HKLM\software\ibm\GSK8`
12. Remove any IBM Spectrum Protect Snapshot user configuration files by entering the following command. Repeat the command for any user accounts that are configured with IBM Spectrum Protect Snapshot:

- a) Change to the following directory:

```
cd %userprofile%\appdata\local\microsoft_corporation
```

Add quotation marks around the directory name if the name contains any spaces. For example:
`cd /d "%userprofile%\appdata\local\microsoft_corporation"`

- b) Enter this command:

```
dir _fmux*
```

- c) Remove each folder that begins with `_fmux`.

Make sure to enclose the folder name in quotation marks (" "). For example:

```
id /s "_FmUx,_Version=4.1.2.0,_C_Path_rusomschqavk3w2upyovnjy1331z5qn3"
```

Chapter 4. Configuring

You can use configuration wizards to configure IBM Spectrum Protect Snapshot, or you can complete the steps manually. For best results, be guided by the step-by-step instructions in the configuration wizards.

About this task

IBM Spectrum Protect Snapshot provides the following wizards to guide your configuration tasks:

Standalone configuration wizard

Use this wizard to configure IBM Spectrum Protect Snapshot to manage snapshot backups as a stand-alone computer. When you select the **Standalone Configuration** option, you configure IBM Spectrum Protect Snapshot to manage snapshots locally, without using an IBM Spectrum Protect server. For stand-alone support, backups are stored locally on the server that is running the backup. The VSS backup is created by using Microsoft Volume Shadow Copy Service. The VSS backup produces an online snapshot (point-in-time consistent copy) of Exchange Server, SQL Server, or custom application and file system data.

IBM Spectrum Protect configuration wizard

Use this wizard to configure IBM Spectrum Protect Snapshot to work with IBM Spectrum Protect. This integration provides data protection and centrally managed, policy-based administration.

When you select the **TSM Configuration** option, IBM Spectrum Protect Snapshot software protects and manages Exchange Server, SQL Server, or custom application and file system data by storing backups locally or on the IBM Spectrum Protect server. With IBM Spectrum Protect, you can also offload your backups to another computer and to move the data to the IBM Spectrum Protect server.

Mailbox Restore Only configuration wizard

Use this wizard to configure IBM Spectrum Protect Snapshot to restore mailboxes from mounted Exchange database EDB files. When you select the **Mailbox Restore Only** configuration option, extra data protection features are not available. This configuration option is ideal when you want to restore mailboxes from only .EDB files and you do not want to use the additional IBM Spectrum Protect Snapshot software functions. The functions that are available with this configuration option are included in the other configuration options.

Specifying configuration parameters for IBM Spectrum Protect

After IBM Spectrum Protect Snapshot for Windows is registered to IBM Spectrum Protect, you must configure the node name, password, the communications method, and the appropriate parameters to connect to the IBM Spectrum Protect server.

Before you begin

When you manually set configuration parameters for IBM Spectrum Protect, ensure that the IBM Spectrum Protect Snapshot options file (`dsm.opt`) and the backup-archive client options file (also `dsm.opt`) specify the same IBM Spectrum Protect server.

About this task

You can use the configuration wizard to set the configuration parameters. To manually configure the parameters, complete the following steps.

Procedure

1. If you are running IBM Spectrum Protect Snapshot on a Microsoft Windows Failover Clustering or Veritas Cluster Server, ensure that the options files on each cluster node are identical by editing the options file. Use a text editor to edit the file.

The `dsm.optoptions` file includes the following parameters, which are necessary for initial configuration:

COMMMethod

Specify the communication protocol to use between the IBM Spectrum Protect Snapshot node and the IBM Spectrum Protect server. Depending on the `commmethod` option that you choose, specify one of the following connectivity parameters for the `commmethod` values.

- For VSS backups, specify the **COMMMethod** option in the IBM Spectrum Protect Snapshot options file.
- For VSS backups, specify the **COMMMethod** option in the backup-archive client options file that is used as the Local DSMAGENT Node. If the environment is configured for VSS offloaded backups, you must also specify the **COMMMethod** option in the backup-archive client options file that is used as the Remote DSMAGENT Node.

NODename

Specify the IBM Spectrum Protect node name that IBM Spectrum Protect uses to identify the system that runs IBM Spectrum Protect Snapshot.

PASSWORDAccess

Specify either the default `generate` value to generate a password automatically, or specify the `prompt` password to respond to a request for a password.

2. Optional: modify the default values for the following parameters:

CLUSTERnode

In the DSMAGENT options files, specify the **CLUSTERnode** option and set it to `no`. For the IBM Spectrum Protect Snapshot options file, specify the **CLUSTERnode** option and set it to `yes`.

HTTPport

Specify the HTTP port. The default value is 1581.

TCPPort

Specify the TCP port.

TCPServeraddress

Specify the TCP server address.

CLUSTERSHARED FOLDER

In the IBM Spectrum Protect Snapshot for SQL Server options file (`tdpsql\dsm.opt`), specify the directory location in which to store encrypted password files. Specify a location that all nodes in the cluster can access. If a failover occurs, the backup-archive client uses this option to determine where the password files are located.

3. For VSS backups that are sent to the IBM Spectrum Protect server, specify the following options to enable features for the data that is sent to the IBM Spectrum Protect server.

When you use these options, you must update the backup-archive client options file that is used as the local DSMAGENT node and the remote DSMAGENT node.

COMPRESSIon

Specify the `compression yes` option if any of the following conditions exist:

- The network adapter has a data overload
- Communications between IBM Spectrum Protect Snapshot and IBM Spectrum Protect server are over a low-bandwidth connection
- Heavy network traffic exists

Specify the `compression no` option if any of the following conditions exist:

- The computer that runs IBM Spectrum Protect Snapshot has a processor overload; the added processor usage might cause issues for other applications that include the server. You can monitor processor and network resource usage with the **Performance Monitor** program that is included with Windows.

- You are not constrained by network bandwidth; you can achieve the best performance by leaving the `compression no` option unchanged and enabling hardware compaction on the tape drive, which also reduces storage requirements.

For VSS backups, specify the **COMPRESSIon** option in the backup-archive client options file that is used as the local DSMAGENT node. If the environment is configured for VSS offloaded backups, specify the **COMPRESSIon** option in the backup-archive client options file that is used as the remote DSMAGENT node.

If you are running offloaded backups, a dedicated (and unique) remote DSM agent node must exist for each local DSM agent node.

DEDUPLICATION

Specify whether the IBM Spectrum Protect API deduplicates data before the data is sent to the IBM Spectrum Protect server. Specify Yes or No. The value applies only if IBM Spectrum Protect allows client-side data deduplication.

When you specify both deduplication and **ENABLELANFree** options, the deduplication option is ignored.

You can enable client-side data deduplication by specifying `DEDUPLICATION YES` in the `dsm.opt` file.

ENABLELANFree

If you run data backup and restore operations in a LAN-free environment, specify **ENABLELANFree yes** in the DSMAGENT (VSS Requestor) options file.

ENABLECLIENTENCRYPTKEY

Specify this option to encrypt databases during backup and restore processing by generating one random encryption key per session.

Restriction: You can back up encrypted VSS databases only to the IBM Spectrum Protect server. You cannot back up encrypted data to an IBM Spectrum Protect Snapshot server.

You can specify DES56 (56 bit), AES128 (128 bit), or AES256 (256 bit). The most secure data encryption method is AES256.

In the options file, you must also specify the databases that you want to encrypt by adding an `include` statement with the `include.encrypt` option.

For VSS backups, specify the encryption options in the backup-archive client options file that is used as the local DSMAGENT node. If the environment is configured for VSS offloaded backups, specify the encryption options in the backup-archive client options file that is used as the remote DSMAGENT node.

If you make changes in the backup-archive client options file, ensure that you restart the IBM Spectrum Protect Client Acceptor Daemon (CAD) service for the Exchange or SQL Server.

What to do next

You can create more IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server options files to point to another IBM Spectrum Protect server. You can create more than one options file, where each file contains different parameters to use with a single IBM Spectrum Protect server.

Related tasks

[“Configuring IBM Spectrum Protect Snapshot for SQL Server clustered environments” on page 70](#)

Depending on the applications that are installed, you can configure IBM Spectrum Protect Snapshot to operate in clustered environments with SQL Server, or file system workloads.

Specifying configuration and options files in non-default locations

The IBM Spectrum Protect Snapshot software uses default configuration and options files. If you want to use non-default configuration and options files, use command-line parameters to specify alternative configuration and option files when you start IBM Spectrum Protect Snapshot.

Before you begin

The information in this procedure does not apply to managing remote IBM Spectrum Protect Snapshot installations.

About this task

MMC that is used for IBM Spectrum Protect Snapshot software is started with the `flashcopymanager.exe` file. The `flashcopymanager.exe` file accepts the following parameters:

```
/mscFilename=filename      # Name of the MMC snap-in control file
/author      # Opens the MMC console in author mode.
```

For example:

```
flashcopymanager.exe parameter1=filename
parameter2=filename ...
```

The `flashcopymanager.exe` file accepts the following parameters to set the configuration files:

```
/FSConfigfile=filename      # File system configuration file
/SQLConfigfile=filename     # SQL configuration file
/EXCCONFIGfile=filename     # Exchange configuration file
/FSOPTfile=filename        # File system OPT file
/SQLOPTfile=filename       # SQL OPT file
/EXCOPTfile=filename       # Exchange OPT file
/FCMCUSTConfigfile=filename # Custom configuration file
/SQLINSTancenames=Instance1,Instance2,... # SQL instances to show in the MMC
```

Procedure

- Start MMC with the parameters by using `flashcopymanager.exe`, as shown in the following example.

```
flashcopymanager.exe /FSConfigfile=newcfg.xml /SQLConfigfile=altsql.cfg
/SQLINSTancenames=mysql1,mysql2
```

You can also start and run multiple instances of MMC concurrently. With the command-line parameters, each instance operates by using a different configuration that is based on the specified configuration and option files.

Setting user preferences

Use the property pages in the **Data Protection Properties** window to customize your IBM Spectrum Protect Snapshot configuration preferences.

Before you begin

The property pages customize preferences such as logging of activity, how languages and information are displayed, and tune performance. The information about the General property page is required to back up data, but the properties are set when you complete the configuration wizard.

When you configure preferences, consider the backup strategy, resource needs, policy settings, and hardware environment of your system.

Procedure

1. In the navigation tree of Microsoft Management Console (MMC), select the Exchange Server, SQL Server, or file system instance for which you want to edit preferences.
2. In the **Action** pane, click **Properties**.
3. Select the property page that you want to view or edit.
The property pages that are available depend on whether your IBM Spectrum Protect Snapshot system is configured for stand-alone support or IBM Spectrum Protect support.
4. Edit the property page and click **OK** to save your changes and close the window.

What to do next

Tip: You can also view or edit properties for the dashboard and MMC. To open the properties window, click **Dashboard** in the navigation tree, and click **Properties** in the **Actions** pane.

Data Protection properties

Use property pages to customize your configuration preferences.

The available property pages for a workload vary depending on whether it is configured for the stand-alone environment or the IBM Spectrum Protect environment.

You can view or edit property pages by selecting a workload from the **Protect and Recover Data** node in the navigation tree of Microsoft Management Console (MMC), and clicking **Properties** in the **Actions** pane.

Server Information

Use the **Server Information** property page to obtain information about the server that provides backup services.

The fields that display depends on whether the product is configured for a stand-alone snapshot environment or for an IBM Spectrum Protect environment.

Note: References to the stand-alone snapshot environment are specific to IBM Spectrum Protect Snapshot.

Node name

Specifies the name that is used to identify the client node for stand-alone backup operations or backup operations to IBM Spectrum Protect server.

TSM API version

Specifies the version of the IBM Spectrum Protect application programming interface (API).

Server name

For backups to IBM Spectrum Protect, specifies the name of the IBM Spectrum Protect server that you are connected to.

For a stand-alone configuration, `Virtual Server` is displayed.

Server Network Host name

Specifies the network host name for the IBM Spectrum Protect server.

For a stand-alone configuration, **FLASHCOPYMANAGER** is displayed.

Server type

For backups to IBM Spectrum Protect, specifies the type of operating system of the IBM Spectrum Protect server.

For a stand-alone configuration, `Virtual Platform` is displayed.

Server version

Specifies the version of the IBM Spectrum Protect server.

Compression mode

Specifies whether compression is used during backup operations to the IBM Spectrum Protect server. The possible values are Yes, No, and Client Determined.

Domain name

Specifies the policy domain that the node belongs to. A policy domain contains one or more policy sets.

For SQL systems, the domain name, policy set, and management class are listed for the Data Protection or AlwaysOn node.

For Exchange systems, the domain name, policy set, and management class are listed for the Data Protection node. To get these parameters for the DAG node, log on to the IBM Spectrum Protect server or contact your IBM Spectrum Protect server administrator.

Active Policy Set

Specifies the policy set that is active for the policy domain. A policy set contains one or more management class definitions.

Default Management Class

The default policy or management class that contains attributes. These attributes determine how long backup versions are stored, where backup versions are stored, and how many backup versions are retained.

Server Password

Use the **Server Password** property page to change the password for the Data Protection node that you use to access the IBM Spectrum Protect server. This property page applies only to IBM Spectrum Protect configurations.

The following fields are displayed in the property page:

Old password

Type the IBM Spectrum Protect password that you want to change.

New password

Type a new password. Follow the IBM Spectrum Protect server password policy rules.

Confirm new password

Type the new password again. Click **OK** to save your changes.

Policy Management

Use the **Policy Management** property page to add or update a backup policy, which controls how different backup versions are retained on local shadow volumes on stand-alone snapshot configurations.

Backup retention on local shadow volumes is determined by version and time-based policies. Ensure that sufficient local storage space is available on local shadow volumes for your VSS backup. The amount of storage space that is required depends on the VSS Provider that you use.

The following fields are displayed in the property page:

Policy

Specify the unique name of a backup policy for the stand-alone configuration.

Number of Snapshots to keep

Specify the number of backup versions to retain on local shadow volumes. Enter a value from 1 to 9999. Type NL to retain as many backup versions as permitted by available storage space. The default value is 2.

This parameter does not apply to incremental backup versions of Exchange Server data. Incremental backups do not participate in expirations because of version limit because there is never more than one version of an incremental backup object. There is only one version of an incremental backup object because incremental backups are always uniquely named.

Days to keep a Snapshot

Specify the number of days to retain backup versions on local shadow volumes. Enter a value from 0 to 9999. Type NL to retain as many backup versions as the available storage space allows. When the value is set to 0, snapshots are kept for the current day. The default value is 30.

VSS Policy Binding

Use the **VSS Policy Binding** property page to bind storage snapshots to back up policies or management classes. VSS policies determine how backups are managed and retained.

VSS policy statements are processed from the end to the beginning and processing stops when the first matching statement is reached. To ensure that more specific statements are processed, specify the more general specification before the more specific ones.

The policy statements do not take effect on existing or new backups until the next backup is completed.

Managed Capacity

Use the **Managed Capacity** property page to track the capacity of managed storage.

The information that is provided can assist you with storage capacity planning during activities such as license renewal.

Diagnostics

Use the **Diagnostics** property page to select the type of tracing to run on various components of IBM Spectrum Protect Snapshot.

When you encounter a problem, open the **Diagnostics** property page. Select the diagnostic mode that you want to use by clicking **Normal**, **Complete**, or **Custom**. Then, click **Begin** to start the trace. Close the property page. Re-create the problem, open the **Diagnostics** property page, and click **End** to stop the tracing and collect the data.

If you are using this property page from the **Dashboard** property sheet, you can run trace only for Microsoft Management Console (MMC).

Diagnostic modes

The following diagnostic mode is available in the **Diagnostics** property page from the **Dashboard** property sheet:

MMC - use this mode to set tracing for MMC only. Only MMC tracing can be completed with this mode.

The following diagnostic modes are available in the **Diagnostics** property page in the workload property sheets. The type of tracing that is enabled for each mode is listed in the table. Specific trace flags, and guidance on when to use each mode is also listed.

Mode	Components traced along with trace flags used	When to use
Normal	MMC, DP (service), API (service, api_detail)	If you are completing legacy operations, you can use this mode because it results in small output size
Complete	MMC, DP (service), API (service, api_detail), Agent (service)	Use for VSS operations, results in large output size
Custom	Any combination	Use if specific flags are needed

Normal

Click **Normal** to collect trace and log files for legacy operations. Not applicable for IBM Spectrum Protect Snapshot for Microsoft Exchange Server.

Complete

Click **Complete** to collect trace and log files for VSS operations.

Custom

Click **Custom**, then click the check mark icon to select the trace and log files that you want to collect. Use this mode only if specific trace flags are required.

Enable snap-in tracing

Select this box to enable tracing of MMC. Click **Review** to view the trace file.

Set Default Trace Flags

Click **Set Default Trace Flags** to set the most commonly requested trace flags.

Enable Data Protection tracing

Select this box to enable tracing of IBM Spectrum Protect Snapshot for Microsoft Exchange Server, IBM Spectrum Protect Snapshot for Microsoft SQL Server, and file system and custom application operations. Click **Review** to view the trace file. Add or update trace flags in the field.

Enable DSM Agent tracing

Select this box to enable tracing for the IBM Spectrum Protect client node. You must restart the client acceptor service before you start the trace. Click **Review** to view the trace file. Add or update trace flags in the field.

Enable API tracing

Select this box to enable tracing for the IBM Spectrum Protect API. Click **Review** to view the trace file. Add or update trace flags in the field.

Email

Select diagnostic files and click **Email** to send a diagnostic email to an IBM service representative with the selected files attached. You must configure your email information before you can send the data to an IBM service representative. To configure your email information, go to the **Dashboard** and click **Properties**. Then, click **Email** to open the email property page.

Screenshot

This function is enabled after you click **Begin**. Click **Screenshot** to open the **Diagnostic Screenshot Tool**. This tool is a modeless dialog that remains open until you close it or click **End** or **Cancel**.

Click **Add New Screenshot** to add a screen capture to the FlashCopyManager\ProblemDetermination folder. The screen capture can be selected with other diagnostic data.

SQL Login

Use this property page to set preferences for logging on to the Microsoft SQL Server. This property page is available only for the **SQL Server** workload.

Use Windows Authentication

Select this option to use a trusted connection and allow Microsoft Windows to authenticate the logon.

Use SQL Server Authentication

Select this option to use SQL user ID security. With this type of security, you must enter the logon ID and the password to log on to the Microsoft SQL Server.

User name

Specifies the SQL user ID.

Password

Specifies the password to log on to the Microsoft SQL Server.

General (SQL Server)

Use **General (SQL)** property page to specify general preferences for the **SQL Server** workload. This property page applies if the product is configured to back up data to stand-alone storage or IBM Spectrum Protect.

SQL Server

Specify the unique name that identifies the SQL Server instance.

From Server

Specify the SQL Server backups that you want to use for the restore. By default, this field displays the same name for the **SQL Server**.

Wait for tape mounts for backup or restore

Select this box when you want IBM Spectrum Protect Snapshot for Microsoft SQL Server to wait for tape media to be mounted for backup and restore operations. This setting is applicable when the IBM Spectrum Protect server is configured to store the backup data on tape media. With backup data on removable media, during backup and restore operations, a wait period occurs during storage volume mounts. If a wait occurs, this setting specifies whether IBM Spectrum Protect Snapshot for Microsoft SQL Server waits for the media mount or stop the current operation. By default, this option is not selected.

Use VSS backups as the default backup method

Select this box to set VSS backups as the default backup method. Ensure that the **Local DSMAGENT Node name** field is specified in the **VSS Backup** property page. Backups can be restored only by using VSS.

Compress backup by using SQL Server compression

Select this box to enable SQL Server compression during legacy backup operations. This check box is available only if you are running Microsoft SQL Server 2008 or later versions.

Compute SQL Server checksum for legacy backup

When selected, this option is written to the IBM Spectrum Protect Snapshot for Microsoft SQL Server preferences file (tdpsql.cfg), and can be applied to all legacy backups. If you clear the check box, you ensure that the integrity check does not apply to any legacy database backup.

Estimate % change for differential backup

Specify the value for the estimated change to database pages for differential backups. This estimate is used by IBM Spectrum Protect Snapshot for Microsoft SQL Server to determine whether enough storage space is available for the backup. The default value is 20. This value becomes the default value for all differential backups.

This field applies only to IBM Spectrum Protect Snapshot for Microsoft SQL Server legacy backups.

General (Exchange Server)

Use the **General (Exchange)** property page to specify general preferences for the **Exchange Server** workload. This property page applies only if your workload is configured to back up data to IBM Spectrum Protect.

Temporary log restore path

Specify the default temporary path to use when you restore logs and patch files. For best performance, specify a path that is on a different physical device than the current active logger. If you do not enter a path, the default is the value of the TEMP environment variable. When you run a full restore, copy restore, or database copy restore, all log files that are in the specified path are erased.

Back up DAG databases to common node

Specify the node name that you want to use to back up databases from a Database Availability Group (DAG). With this setting, all active and passive copies of the databases are backed up to the same file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they were backed up from.

When you use this setting, IBM Spectrum Protect applies the same policy across all DAG members, regardless of which DAG member ran the backup.

Temporary database restore path

Specify the directory where the database files that are being restored are temporarily located. Ensure that the directory provides enough space to store the entire mailbox database file. If you do not specify a directory, the database files are restored into a directory that is specified by the TEMP environment variable. This option is only available for mailbox restore operations.

Alias of temporary mailbox

Specifies the alias of a mailbox to use as a temporary storage location during mailbox restore operations. The temporary mailbox is used during restore operations of mailboxes that were deleted,

re-created, or moved since the time of the backup. By default, the mailbox restore operation uses the administrator user's mailbox as a temporary storage location.

Exchange Client Access Server

Specify the name of the Client Access Server (CAS) that you want to use. This field is available only for Microsoft Exchange Server 2013 or later versions.

By default, IBM Spectrum Protect Snapshot uses the local server as the CAS if the local server has the CAS role installed. The CAS that is defined by the logon user mailbox database is used if the local server does not have the CAS role installed.

You can find the name of the current CAS, which is defined by the current logon user mailbox database, by running this Exchange Management Shell command:

```
Get-MailboxDatabase -Identity <logon user mailbox database> |  
select RpcClientAccessServer
```

To use a different CAS, you can define the CAS to be used.

Restore mail messages as unread

Select this check box to specify that restored mail messages are marked as unread.

Backup mailbox history

Select this check box if you are using mailbox restore operations and you want the mailbox history to be backed up.

Tip: If you do not intend to run mailbox restore operations, clear this check box. This action can improve backup performance.

Pre/Post Snapshot

Use this property page to specify presnapshot and postsnapshot commands. This property page applies only to custom applications in the **File System** workload.

Pre-Snapshot Command

Specify the name of the command script that is used to quiesce custom applications that use the file system before the snapshot is created. You must specify the fully qualified path name for the command script.

Post-Snapshot Command

Specify the name of the command script that is used to restart custom applications that use the file system after the snapshot is created. You must specify the fully qualified path name for the command script.

All batch scripts must include an exit statement with the following value:

```
exit error_code
```

Logging

Use the **Logging** property page to specify activity log preferences.

Log File Name

Specifies the name of the file in which activities are logged.

Enable pruning

Specifies that older entries from the log are to automatically be deleted. By default, log pruning is activated and completed daily.

Number of days to keep old entries

Specifies the number of days to keep old entries in the log before they are pruned. By default, 60 days of log entries are saved in the pruning process.

Prune now

Click this option to delete older entries from the IBM Spectrum Protect Snapshot activity log when a command runs.

Regional

Use the **Regional** property page to set preferences that affect how languages and information are displayed and logged.

Regional and Language options

Select this option to set preferences for Microsoft Management Console (MMC). MMC uses the same regional settings as the Windows system.

Language

Select the language to use for log files and the command-line interface.

Date Format

Select a date format to use for log files and the command-line interface. The available choices represent several ways to place the month (*mm*), day (*dd*), year (*yyyy*), and period of day (*a.m.* or *p.m.*). The default date format is *mm/dd/yyyy*.

Time Format

Select a time format to use for log files and the command-line interface. The available choices represent several ways to place the hour (*hh*), minutes (*mm*), and seconds (*ss*). The default time format is *hh:mm:ss*.

Number Format

Select a number format to use for log files and the command-line interface. The available choices represent several ways to place the decimal, comma, and spaces. The default number format is *xxx,xxx.dd*.

Match MMC Language

Select this option to change MMC regional settings to match the system's regional and language options. By selecting this option, the number, date, and time formats are matched to the default formats of the selected language.

VSS Options

Use the **VSS Options** property page to configure preferences that are used during VSS backup and restore operations.

Default Backup Destination

Select the default storage location for data backups.

Tip: You must have the IBM Spectrum Protect Snapshot license to use the IBM Spectrum Protect software. If you have only the Data Protection license, only the IBM Spectrum Protect option is enabled.

You can select from the following storage locations:

IBM Spectrum Protect server

The backup is stored only on IBM Spectrum Protect server storage. Select this option for workloads that are configured with IBM Spectrum Protect server. For Exchange Server and SQL Server, IBM Spectrum Protect server is the default backup destination.

Local

The backup is stored only on a local disk. For custom application and file system data, a local backup is the default backup destination.

Both

The backup is stored on IBM Spectrum Protect server storage and on a local disk. Select this option for workloads that are configured with IBM Spectrum Protect server.

For IBM Spectrum Protect configurations, the backups can be stored on a local disk, but managed on the IBM Spectrum Protect server. The IBM Spectrum Protect server maintains the metadata, that is, the information about where the local snapshot is stored.

Local DSMAGENT Node name

Specify the node name for the DSM Agent node of the local client system that creates the VSS backups.

Remote DSMAGENT Node name

Specify the node name of the system that moves the VSS data to IBM Spectrum Protect server storage during offloaded backups. If you do not use offloaded backups, you can leave this field blank.

If you are running offloaded backups, a dedicated (and unique) remote DSM agent node must exist for each local DSM agent node.

Import VSS snapshots only when needed

Select the check box to have IBM Spectrum Protect Snapshot import VSS snapshots to the Windows system where the snapshots are created. The check box is selected by default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required.

Ensure that the check box is selected if you want to take the following actions:

- Run instant restore operations on some IBM and non-IBM storage systems by enabling the storage system to create transportable snapshots during backups
- Import the VSS snapshots to a local server
- Keep more than 100 backup versions
- Extend the number of LUNs that the server can use, for example, in a VMWare environment

Tip: If you work in a VMware environment and want to use VMware vMotion, ensure that the LUNs are correctly zoned to the ESX hosts. The import process maps the VSS snapshot to the ESX host where the Windows virtual machine is running.

Clear the check box if you do not want to create transportable VSS snapshots during backup processing and automatically import the snapshot to storage systems after the backup is completed.

During Instant Restore, automatically stop and restart necessary Microsoft Exchange services

When this option is selected, during instant restore operations, the following Microsoft Exchange services are, as necessary, automatically stopped and restarted:

- (DAG environments only) Exchange Replication Service
- (Exchange 2013, 2016, and 2019 only) Exchange Search Host Controller Service

Mount read only

Select the check box to specify that backups are to be mounted as read-only VSS snapshots by default. However, at mount time, you can override this value and do a read/write mount. If you change the default, the corresponding update is made in your configuration file automatically.

Mount read/write (modifies backup, applies to COPY backups only)

Select the check box to specify that backups are to be mounted as read/write VSS snapshots by default. You can mount only COPY backups as read/write and after mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations. However, at mount time, you can override this value and do a read-only mount.

Mount read/write (without modifying backup)

Select this check box to specify that backups are to be mounted as read/write copies of the backup by default. With this option, you can mount both FULL and COPY backup types as read/write. After mounting, the original backup is not modified and you can use it again in future database restore operations. However, at mount time you can override this value and do a read-only mount.

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which requires IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which requires IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate extra target volumes on your SVC storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted is needed for each concurrent read/write mount of that volume.

Capacity Licensing

Use the **Capacity Licensing** property page to generate capacity metrics reports for users who are licensed under the front-end or back-end capacity licensing model.

The report, in an XML file format, is generated automatically at the end of a backup operation. You can specify the frequency at which the report is generated and the location into which it is created by using the following options on the **Capacity Licensing** property pane.

XML creation interval

Specify the frequency in days at which the report is automatically generated at the end of a backup operation. Valid values are 1 - 365 and the default value is 7 days.

XML creation location

Specify the directory path into which the report is created. The location you specify is a network shared folder that can be accessed by the central reporting tool.

You must specify a creation location. If you do not, the report is not generated.

Custom Settings

Use the **Custom Settings** property page to set your filtering options and control the amount of information that is returned from the server.

Select **Show Refresh Options** in the toolbar in the **Recover** view. This property page is available only with SQL Server and Exchange Server workloads.

In environments where thousands or millions of backup objects are stored on the IBM Spectrum Protect server, it can be helpful to disable the automatic refresh mode. You can click **Refresh Options** and use the toolbar to switch between manual and automatic refresh mode.

Automatic and manual refresh modes differ in the following ways:

- In automatic refresh mode, a view automatically refreshes the first time that you select it. If there are thousands or millions of objects on the server, the refresh can take a long time to complete.
- In manual refresh mode, the view is not automatically refreshed. A name filter is available on the **Refresh Options** toolbar that you can use to narrow down the number of objects selected. After you enter a name pattern, you can click **Refresh**. By using manual refresh mode and limiting your query by using filters, you can reduce the amount of information that is returned from the server. Reducing the amount of information that is returned from the server can improve query and restore performance.

To help you to filter information, you can also specify an asterisk (*) as a wildcard character in the name pattern.

Performance

Use this property page to set preferences that affect performance for legacy backups. This property does not apply to Data Protection for Exchange Server.

DP Buffers

Specifies a number in the range 2 - 8 that specifies the number of communication data buffers that IBM Spectrum Protect Snapshot for Microsoft SQL Server uses when it transfers data to the IBM Spectrum Protect server. Each buffer has the size that is specified by the **DP Buffer Size** option. This option applies only to legacy backups.

DP Buffer size

Specifies the size of the buffers that are used by IBM Spectrum Protect Snapshot for Microsoft SQL Server to transfer data to the IBM Spectrum Protect server. This option applies only to legacy backups. Specify a value in the range 64 - 8192.

SQL Buffers

Specifies the number of communication data buffers that IBM Spectrum Protect Snapshot for Microsoft SQL Server uses when it transfers data between the SQL Server and IBM Spectrum Protect Snapshot for Microsoft SQL Server. Each buffer has the size that is specified in the **SQL Buffer Size** option. This option applies to legacy backups only. Specify a value from 0 - 999.

SQL Buffer size

Specifies the size of the buffers that are used by IBM Spectrum Protect Snapshot for Microsoft SQL Server to transfer data from the SQL Server to IBM Spectrum Protect Snapshot for Microsoft SQL Server. This option applies only to legacy backups. Specify a value in the range 64 - 4096.

Stripes

Specifies the number of data stripes to use in a legacy backup or legacy restore operation. Specify a value in the range 1 - 64. The default value is 1. This option applies to legacy backup and restore operations only.

When you use a multiple stripes number for legacy backups, and setting the **Verify Only** parameter to **Yes** to restore the legacy backup, the number of stripes for legacy restore must be equal or greater than the number of stripes for the legacy backup.

AlwaysOn Node

All availability databases in an availability group are backed up under this AlwaysOn node.

When you configure IBM Spectrum Protect Snapshot in a SQL Server 2012 environment, the AlwaysOn node name is a required parameter. The AlwaysOn node name can be changed at any time. To change the parameter, use the configuration wizard. From the AlwaysOn Node property page, view the parameter that is set.

AlwaysOn node name

The AlwaysOn node name that is used to back up availability databases is displayed. The databases that are not in an availability group are backed up under the standard node name unless you select the **Use AlwaysOn node name for all databases** check box.

You cannot change the node name from this property page. To change the node name, use the configuration wizard.

Use AlwaysOn node name for all databases

Select this check box to specify that the AlwaysOn node is the default node for backing up all availability and standard databases. This option can be used to change your database backups from the standard node to an AlwaysOn node. By selecting this check box, you can back up all standard and availability databases under a single node to help you manage your database backups more easily.

AlwaysOn Preferences

Use this property page to configure the backup preference settings for scheduled backups of availability groups and availability replicas in an SQL Server 2012 environment.

The settings are intended for scheduling backups of availability groups and availability replicas, and not for interactive backups with the Microsoft Management Console (MMC).

The following settings are available in this property page:

Availability group

Select an AlwaysOn Availability Group for which you want to set schedule backup preferences.

Primary replica

Displays the primary replica for the selected availability group.

Preferred replica

Backup commands are scheduled on all nodes in an availability group. When the schedules run, the backup occurs only on the preferred replica. Other replicas receive, at run time, a warning message. The backup is skipped. You can make the following selections:

- Select **Prefer Secondary replica** if you want scheduled backups to occur on a secondary replica, if it is available. Otherwise, use the primary replica for the scheduled backup.
- Select **Secondary only** if you want scheduled backups to occur only on a secondary replica.
- Select **Primary** if you want scheduled backups to occur only on the primary replica.
- Select **Any replica** if you want scheduled backups to occur on any availability replica.

Availability replicas

For each availability replica in this list box, specify whether it is a candidate for running scheduled backups by specifying the backup priority for that replica. A value of 1 has the lowest priority, and a value of 100 has the highest priority. A value of 0 indicates that the replica is excluded from schedule backup operations.

Availability databases

Displays the availability databases that are in an availability replica. The synchronization state is also displayed.

MAPI Settings

If you use Exchange Server 2013, use the **MAPI Settings** property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables IBM Spectrum Protect Snapshot to connect to the Exchange Server.

IBM Spectrum Protect Snapshot automatically generates a default value for the registry key. Edit the registry key only if the default value is incorrect.

The values that you enter override the registry key that is in the HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem directory. If you modify the registry incorrectly, the connection to the Exchange Server might fail.

You can use this property page only if you use Exchange Server 2013.

RpcHttpProxyMap_TSM

Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment. By default, the format is:

```
Domain=Proxy Server,RpcHttpAuthenticationMethod,  
RpcAuthenticationMethod,IgnoreSslCert
```

For example:

```
companyname.local=https://exchange.companyname.com,ntlm,ntlm,false
```

where:

- *Domain* value is the domain suffix of the personalized server ID, for example, `companyname.local`. Specify any domain or a substring of a domain, or the asterisk (*) and question mark (?) wildcard characters, for example, `*.companyname.local`.
- *Proxy Server* value is the RPC proxy server that has the Client Access Server (CAS) role. Specify the fully qualified domain name (FQDN) of the RPC proxy server. Precede the FQDN by `http://` for an HTTP connection, or `https://` for an HTTPS connection. For example, `https://exchange.companyname.com`
- *RpcHttpAuthenticationMethod* value is the method that is used to authenticate RPC-over-HTTP connections. Specify NTLM, Basic, Negotiate, or WinNT.
- *RpcAuthenticationMethod* value is the method that is used to authenticate RPC-over-TCP connections. Specify NTLM, Negotiate, WinNT, Anonymous, or None.
- *IgnoreSslCert* value indicates whether the Exchange Server validates SSL certificates. For the Exchange Server to ignore invalid certificates, specify `False`.

Domain

Change the domain name to reflect the correct domain if for example, you have multiple domains, or the default domain value is incorrect. To match all domains, enter the asterisk (*) wildcard character. When you change this domain value, the *Domain* value of the registry key automatically updates in the `RpcHttpProxyMap_TSM` field.

Use HTTPS authentication

Select this check box if RPC-over-HTTPS is enabled for the Exchange Server that is hosting the MAPI profile. Otherwise, clear this check box to ensure that HTTP authentication is used for RPC-over-HTTP

connections. When you change this authentication value, the *RpcAuthenticationMethod* value of the registry key automatically updates in the *RpcHttpProxyMap_TSM* field.

Related tasks

Ensuring successful MAPI connections

If you use Exchange Server 2013, use the **MAPI Settings** property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables IBM Spectrum Protect Snapshot to connect to the Exchange Server.

Configuring IBM Spectrum Protect Snapshot in a stand-alone configuration

By using the **Standalone Configuration Wizard**, you can configure IBM Spectrum Protect Snapshot to store database snapshots locally, without using an IBM Spectrum Protect server.

About this task

The configuration procedure applies to the following environments:

- IBM Spectrum Protect Snapshot for SQL Server, if the required license is located
- IBM Spectrum Protect Snapshot for Exchange Server, if the required license is located
- File system and custom applications

Procedure

1. To start Microsoft Management Console (MMC), click **Start > All Programs > IBM Spectrum Protect Snapshot > IBM Spectrum Protect Snapshot Management Console**.

If the **Management Console** is not configured for licenses that are identified, a welcome page is displayed. You can select the type of configuration to complete.

2. From the start page, click **Configuration**.

You are automatically prompted to run the configuration wizard. If the configuration wizard prompt is not displayed automatically, expand the following tree nodes: **IBM Spectrum Protect > Dashboard > Manage > Configuration**. Then, select **Wizards**.

3. In the results pane, double-click **Standalone Configuration** to open the **Standalone Configuration Wizard**.
4. Follow the wizard instructions to configure stand-alone snapshot support.

- a) In the **Data Protection Selection** page, select the applications that you want to protect.

You can select the **SQL Server**, **Exchange Server**, or **File System** workload.

- b) To view information about the computer, operating system, processor, and physical memory, click **Show System Information**.

- c) Review the information in the **Requirements Check** page. Correct any error or warning messages.

For Exchange Server workload, if you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient role-based access control (RBAC) permissions, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.

- d) Select the **Default** configuration setting.

When you select the **Default** configuration setting, the VSS Requestor is configured in addition to configuring the applications that you selected. The client and agent services are also registered and configured, and a schedule to support historical managed capacity is defined.

If you need more than one instance of the Client Acceptor and Remote Agent services, or if you use the backup-archive client to back up to the IBM Spectrum Protect server, but IBM Spectrum Protect Snapshot is in a stand-alone configuration, click **Custom**. You can use the **Custom** setting

to choose a node name for the Client Acceptor and Remote Agent services, an options file, service names, and the HTTP port. IBM Spectrum Protect Snapshot does not interfere with the existing client operations. If the backup-archive client is not installed and configured to protect the file system, the **Default** setting is easier to use.

When you select **Custom**, more fields are displayed to change the client service configuration. Review the information in the fields and, if necessary, change settings.

- **VSS Requestor node name:** Enter the node name that communicates with the VSS Service to access the Exchange Server, SQL Server, or custom application and file system data. The VSS Requestor node name is also the node name that the Remote Client Agent service uses to communicate with IBM Spectrum Protect Snapshot.
- **VSS Requestor options file name:** Enter the name of the client options file for the VSS Requestor node.
- **Client Acceptor service name:** Specify the name of the service that is used by IBM Spectrum Protect backup-archive client to communicate with IBM Spectrum Protect Snapshot. By default, this service is named the **IBM Spectrum Protect Client Acceptor**.
- **Remote Client Agent service name:** Specify the name of the service that communicates with Windows VSS to run the VSS operations. By default, this service is named the **IBM Spectrum Protect Remote Client Agent**.
- **HTTP Port:** Specify the HTTP port to use for the Client Acceptor service.

You can also delete an existing service by selecting a service in the **Currently installed client services** list and clicking **Remove**. Removal of a service happens instantly. The removal occurs when you click **Remove**.

- e) Click **Show Details** to view a list of individual configuration results.
5. Click **Finish** to complete the wizard.
6. Optional: After you complete the configuration process with the wizard, test VSS snapshots on the system. Click **Run VSS diagnostics when this wizard exits**.



Attention: If the configuration is for space-efficient target volumes for SAN Volume Controller or Storwize V7000, testing VSS snapshots deletes previous backups that are created for the volumes that are selected in the test wizard.

7. To verify that IBM Spectrum Protect Snapshot is correctly configured, select a workload in the **Protect and Recover Data** node in the navigation tree. From the **Automate** view, issue one of the following commands.
. For example, the following CLI commands can be used:

- For file systems and custom applications:

```
fcmcli query component  
fcmcli query config
```

- For SQL Server:

```
tdpsqlc query tdp  
tdpsqlc query fcm  
tdpsqlc query sql
```

- For Exchange Server:

```
tdpexcc query tdp  
tdpexcc query fcm  
tdpexcc query exchange
```

You can use the selection tool to choose to enter either CLI commands or PowerShell cmdlets. You can also view the configuration settings by clicking **Properties** for each configured workload.

What to do next

After you complete the configuration wizard, you can use IBM Spectrum Protect Snapshot to back up and restore data.

Configuring an IBM Spectrum Protect Snapshot remote system in a stand-alone configuration

By using the **Standalone Configuration Wizard**, you can configure a remote system to work in a stand-alone environment.

Before you begin

On the local system, verify the following system requirements:

- Windows 7, Windows 8, Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed
- PowerShell version 3.0 or later is installed, if you are running Windows 7, Windows 8, Windows 2008, or Windows 2008 R2. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default.
- IBM Spectrum Protect Snapshot version 4.1.4 is installed

On the remote system, verify the following system requirements:

- Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed
- Windows PowerShell version 3.0 or later is installed, if you are running Windows 2008, or Windows 2008 R2. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default.
- IBM Spectrum Protect Snapshot version 4.1.4 is installed
- The required workload is configured.

Procedure

1. On the local system, start the IBM Spectrum Protect Snapshot **Management Console**.
2. From the **Management Console**, use **Manage Computers** to add the remote system.
3. In the navigation tree, verify that the remote system is displayed.
4. Click **Manage > Configuration > Wizards**.
5. Select **Standalone Configuration**.
6. On the **Data Protection Selection** page, verify that the following information is entered correctly:
 - The remote computer name in the window title.
 - The correct system information.
7. Select the application to be configured and click **Next**.
8. On the **Requirements Check** page, click **Show Details**.

For Exchange Server workload, if you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient role-based access control (RBAC) permissions, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.
9. On the **Custom Configuration** page, select **Default**.
10. On the **Configuration** page, click **Show Details**.

Verify the progress and status of the configuration.
11. Click **Finish** to complete the wizard.

What to do next

To verify that the configuration is set up correctly, complete the following steps:

1. In the navigation tree, for the remote system, expand **Protect and Recover** and click the application that is configured.
2. Query the components and verify that a successful backup can be completed.

Related concepts

Security requirements for backup and restore operations

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Configuring IBM Spectrum Protect Snapshot to integrate with IBM Spectrum Protect

By using the IBM Spectrum Protect **Configuration Wizard**, you can configure IBM Spectrum Protect Snapshot to protect and manage Exchange Server, SQL Server, or custom application and file system data by storing backups locally or on the IBM Spectrum Protect server.

Before you begin

If you configure the DSM Agent node (the backup-archive client node) manually, ensure that you set the **PASSWORDAccess** option to generate in the `dsm.opt` file for the IBM Spectrum Protect backup-archive client. Also ensure that the stored password for the DSMAGENT Node is valid.

About this task

The configuration procedure applies to the following environments:

- IBM Spectrum Protect Snapshot for SQL Server, if the required license is located
- IBM Spectrum Protect Snapshot for Exchange Server, if the required license is located
- File system and custom applications

Procedure

1. To start Microsoft Management Console (MMC), click **Start > All Programs > IBM Spectrum Protect Snapshot > IBM Spectrum Protect Snapshot Management Console**.
2. From the start page, click **Configuration**.

You are automatically prompted to run the configuration wizard. If the configuration wizard prompt is not displayed automatically, expand the following tree nodes: **IBM Spectrum Protect > Dashboard > Manage > Configuration**. Then, select **Wizards**.

3. In the results pane, double-click **IBM Spectrum Protect Configuration** to open the IBM Spectrum Protect **Configuration Wizard**.
4. Follow the wizard configuration instructions, and click **Next** to move to the next page.
 - a) In the **Data Protection Selection** page, select the applications that you want to protect.
You can select the **SQL Server**, **Exchange Server**, or **File System** workload.
 - b) Review the results of the requirements check and ensure that you address any errors or warnings.
Click **Show Details** to view results.

- If you are configuring an application for which you do not have the necessary license, the license requirement check fails. You must either return to the **Data Protection Selection** page and clear the selected application to proceed with the configuration, or obtain the necessary license.
- For Exchange Server workload, if you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient role-based access control (RBAC) permissions, you can automatically

add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.

- c) In the IBM Spectrum Protect **Node Names** page, specify the IBM Spectrum Protect node names, which exist on the same system, to use for the applications that you want to protect.

<i>Table 10. Field entry in the IBM Spectrum Protect Node Names page</i>	
Field	Action
VSS Requestor	Enter the node name that communicates with the VSS service to access the Exchange Server, SQL Server, or custom application and file system data.
Data Protection for SQL	Enter the node name where the Data Protection application is installed. This name is the node name that is used to store the IBM Spectrum Protect Snapshot for SQL Server backups. Tip: If you do not need a VSS configuration for your SQL Server, you can skip the configuration. Click Do not configure DP SQL VSS support .
AlwaysOn Node	Enter a node name if you are configuring IBM Spectrum Protect Snapshot with SQL Server 2012 and later versions. This name is the node name that is used to back up the availability databases in an AlwaysOn Availability Group. By default, the Windows Failover Cluster name is used.
Data Center Node	Enter the data center node name if the IBM Spectrum Protect for Virtual Environments Recovery Agent license is available. The data center node is the virtual node that maps to a data center.
Data Protection for Exchange	Enter the node name where the Data Protection application is installed. This name is the node name that is used to store the IBM Spectrum Protect Snapshot for Exchange Server backups. If you configure the DAG Node on this wizard page, Exchange Server DAG database backups are stored under the DAG node.

<i>Table 10. Field entry in the IBM Spectrum Protect Node Names page (continued)</i>	
Field	Action
DAG Node	<p>Enter the node name that you want to use to back up databases in an Exchange Server DAG. With this setting, backups from all DAG members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server.</p> <p>The database copies are managed as a single entity, regardless of which DAG member they were backed up from. This setting can prevent IBM Spectrum Protect Snapshot from making too many backups of the same database.</p> <p>Ensure that you configure all of your DAG members that have copies of the same database to all use the same DAG node. On the IBM Spectrum Protect server, ensure that you register the DAG node name. All of the DAG member nodes (the Data Protection nodes) must be granted <i>proxynode</i> authority to run backups on behalf of the DAG node. All of the DSM Agent nodes (the backup-archive client nodes) must also be granted <i>proxynode</i> authority. If you do not want to manually update these properties, you can use the configuration wizard to set the properties on the IBM Spectrum Protect server.</p>
Files System and Custom Configuration	<p>Enter the node name that you want to use to back up custom application and file system data.</p>

Create a node name that can help you distinguish the type of backup that is run. For example, if your host name is *MALTA*, you can name the VSS Requestor node name *MALTA*, and you can create a Data Protection node name that is called *MALTA_EXC* or *MALTA_SQL*. For an SQL Server configuration, the AlwaysOn node name does not have to be related to the VSS Requestor or the IBM Spectrum Protect Snapshot for SQL Server node name. For example, you can name it *TSM_ALWAYSON*. For an Exchange Server configuration, the DAG node name does not have to be related to the VSS Requestor or the IBM Spectrum Protect Snapshot for Exchange Server node name. For example, you can name it *TSMDAG*.

- d) Enter information for the IBM Spectrum Protect server that you are connecting to and click **Next** to continue.

<i>Table 11. Field entry in the IBM Spectrum Protect Node Names page</i>	
Field	Action
IBM Spectrum Protect Server Address	<p>Enter the TCP/IP domain name or a numeric IP address for the IBM Spectrum Protect server that contains the backups. Obtain this information from your IBM Spectrum Protect server administrator.</p>

Table 11. Field entry in the IBM Spectrum Protect Node Names page (continued)	
Field	Action
IBM Spectrum Protect Server Port	Enter the port number for the IBM Spectrum Protect server that contains the backups. Obtain this information from your IBM Spectrum Protect administrator.

Specify whether to have the wizard to configure the IBM Spectrum Protect server for you by generating a configuration macro file.

If you click **No**, the macro file is available at the final page of the wizard so that it can be provided to the IBM Spectrum Protect administrator as an example of one way to configure the IBM Spectrum Protect server to support application data protection.

If you click **Yes**, the wizard starts the macro during the **Configuration** step in the wizard. Review the macro file and update it if needed.

After you click **Yes**, enter the following information in the appropriate field:

- The name of the IBM Spectrum Protect administrator account.
- The password for the IBM Spectrum Protect administrator.
- Click **Test Communications** if you want to test your connection with the IBM Spectrum Protect server. This option is not available until the VSS Requestor is installed.
- Click **Review/Edit** to review or update the IBM Spectrum Protect macro file. Alternatively, you can review the macro file and directly run the commands on the IBM Spectrum Protect server.

e) Select the **Default** configuration setting.

When you select the **Default** configuration setting, the VSS Requestor is configured in addition to configuring the applications that you selected. The client and agent services are also registered and configured, and a schedule to support historical managed capacity is defined.

f) Review the results of the configuration process. Click **Show Details** to view a list of individual configuration results.

5. Click **Finish** to complete the wizard.

6. Optional: After you complete the configuration process with the wizard, test VSS snapshots on the system. Click **Run VSS diagnostics when this wizard exits**.



Attention: If the configuration is for space-efficient target volumes for SAN Volume Controller or Storwize V7000, testing VSS snapshots deletes previous backups that are created for the volumes that are selected in the test wizard.

7. To verify that IBM Spectrum Protect Snapshot is correctly configured, select a workload in the **Protect and Recover Data** node in the navigation tree. From the **Automate** view, issue one of the following commands.

- For file systems and custom applications:

```
fccli query component
fccli query config
```

- For SQL Server:

```
tdpsqlc query tdp
tdpsqlc query tsm
tdpsqlc query sql
```

- For Exchange Server:

```
tdpexcc query tdp
tdpexcc query tsm
tdpexcc query exchange
```


You can use the selection tool to choose to enter either CLI commands or PowerShell cmdlets. You can also view the configuration settings by clicking **Properties** for each configured workload.

What to do next

After you complete the configuration wizard, you can use IBM Spectrum Protect Snapshot to back up and restore data.

Configuring an IBM Spectrum Protect Snapshot remote system to integrate with IBM Spectrum Protect

By using the **TSM Configuration Wizard**, you can configure a remote system to communicate with an IBM Spectrum Protect server.

Before you begin

On the local system, verify the following system requirements:

- Windows 7, Windows 8, Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed
- PowerShell version 3.0 or later is installed, if you are running Windows 7, Windows 8, Windows 2008, or Windows 2008 R2. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default.
- IBM Spectrum Protect Snapshot version 4.1.4 is installed

On the remote system, verify the following system requirements:

- Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed
- Windows PowerShell version 3.0 or later is installed, if you are running Windows 2008, or Windows 2008 R2. On Windows 2012 and later versions, PowerShell version 4.0 is installed by default.
- IBM Spectrum Protect Snapshot version 4.1.4 is installed
- The required workload is configured.

Procedure

1. On the local system, start Microsoft Management Console (MMC).
2. From MMC, use **Manage Computers** to add the remote system.
3. In the navigation tree, verify that the remote system is displayed.
4. Click **Manage > Configuration > Wizards**.
5. Select **TSM Configuration**.
6. On the **Data Protection Selection** page, verify that the following information is entered correctly:
 - The remote computer name in the window title.
 - The correct system information.
7. Select the application to be configured and click **Next**.
8. For Exchange or SQL Server, the license check might fail. If the test fails, provide the file path and name for the location on the remote server.
9. On the **TSM Node Names** page, verify that the following information is entered correctly:
 - VSS Requestor
 - The Data Protection or file system name, depending on the application that is configured

For systems with a Database Availability Group (DAG) or an AlwaysOn Availability Group, the corresponding DAG node or AlwaysOn node is detected.
10. On the **TSM Server Settings** page, type the server name and port number.
11. For the **Would you like this wizard to configure your TSM server?** question, select **Yes**.
12. Click **Review / Edit**. If the domain is not entered correctly, update the information. Click **OK**.
13. On the **Custom Configuration** page, select **Default**.

14. On the **Configuration** page, click **Show Details**.
Verify the progress and status of the configuration.
15. Click **Finish** to complete the wizard.

What to do next

To verify that the configuration is set up correctly, complete the following steps:

1. In the navigation tree, for the remote system, expand **Protect and Recover** and click the application that is configured.
2. Open the **Properties** and click **Server Information**. Verify that the correct information is displayed.
3. Query the components and verify that a successful backup can be completed.

Related concepts

[Security requirements for backup and restore operations](#)

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Configuring IBM Spectrum Protect Snapshot to restore mailboxes from mounted Exchange Server database files

The **Mailbox Restore Only** configuration wizard is useful when you do not have to configure additional IBM Spectrum Protect Snapshot software functions.

Procedure

1. To start Microsoft Management Console (MMC), click **Start > All Programs > IBM Spectrum Protect Snapshot > IBM Spectrum Protect Snapshot Management Console**.
2. From the start page, click **Configuration**.
You are automatically prompted to run the configuration wizard. If the configuration wizard prompt is not displayed automatically, expand the following tree nodes: **IBM Spectrum Protect > Dashboard > Manage > Configuration**. Then, select **Wizards**.
3. In the results pane, double-click **Mailbox Restore Only** to open the **Mailbox Restore Only Configuration** wizard.
4. Follow the wizard configuration instructions. Click **Show Details** to view a list of individual requirement results.

Review the results of the requirements check and address any errors or warnings.

For Exchange Server workload, if you do not have all the user roles that are required for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange Server roles. If you are a member of the Exchange Organization Management group and have sufficient Role Based Access Control (RBAC) permissions, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group and have insufficient RBAC permissions, you must manually add the missing roles.

5. Click **Finish** to complete the wizard.

Related concepts

[Security requirements for backup and restore operations](#)

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Configuring node definitions

Although IBM Spectrum Protect Snapshot can automatically configure node definitions, you can also manually configure node names for IBM Spectrum Protect Snapshot. You can also configure the system that runs offloaded backups.

Proxy node definitions for VSS backups

Use the IBM Spectrum Protect backup-archive client to implement VSS backup operations. You must use two node names for each configured snapshot agent; one for the backup-archive client and the other for the IBM Spectrum Protect Snapshot agent (for Exchange, SQL or Custom Applications).

As part of the configuration procedure, a proxy relationship is defined for these node names. By default, this proxy relationship is defined when you run the configuration wizard. You can manually complete the configuration.

The proxy relationship allows node names to process operations on behalf of another node name. When you register these nodes to the IBM Spectrum Protect server for VSS operations, you must specify the IBM Spectrum Protect `USerid=<node name>` parameter.

Two types of node names are defined in proxy node relationships:

- *Target node*: A node name that controls data backup and restore operations and also owns the data on the IBM Spectrum Protect server. This node name is specified in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server and IBM Spectrum Protect Snapshot for Microsoft SQL Server `dsm.opt` file.
- *Agent node*: A node name that processes operations on behalf of a target node. This node name is specified in the backup-archive client `dsm.opt` file.

To define these nodes, enter the backup-archive client **grant proxy** command. For example:

```
GRANT PROXY TARGET=dpexc_node_name AGENT=dsmagent_node_name
```

Required node names for basic VSS operations

VSS operations require specific node name settings.

To process basic VSS operations, you must have one target node and one agent node.

Table 12. Required node names for basic VSS operations

Proxy node type	Node name	Where to specify
Target node	The IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server node name	Use the <code>nodename</code> option in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server options file (<code>dsm.opt</code>)
Agent node	The Local DSMAGENT Node name that must match the backup-archive client node name	Use the localsmagentnode parameter in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file (<code>tdpexc.cfg</code>) or (<code>tdpsql.cfg</code>)

Note: For basic VSS operations, the agent node and target node are on the same system.

Required node names for basic VSS offloaded backups

VSS offloaded backups require specific node name settings.

To complete VSS offloaded backups, you must have one target node and two agent nodes:

Table 13. Required node names for basic VSS offloaded backups

Proxy node type	Node name	Where to specify
Target node	IBM Spectrum Protect Snapshot for Microsoft Exchange Server IBM Spectrum Protect Snapshot for Microsoft SQL Server node name	Use the nodename option in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server options file (<code>dsm.opt</code>)
Agent node	Local DSMAGENT Node	Use the localdsmagentnode parameter in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file (<code>tdpexc.cfg</code>) or (<code>tdpsql.cfg</code>)
Agent node	Remote DSMAGENT Node	Use the remotedsmagentnode parameter in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file <code>tdpexc.cfg</code> or (<code>tdpsql.cfg</code>)

Target node

This node name is where IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server is installed. This node name (specified with the **nodename** option in the `dsm.opt` file) is referred to as the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server node name.

Agent node - Local DSMAGENT Node

This node name is where the backup-archive client and VSS provider are installed. This node is responsible for processing the VSS operations because IBM Spectrum Protect Snapshot for Microsoft Exchange Server does not process any direct VSS operations.

This node name is referred to as the Local DSMAGENT Node and is specified with the **localdsmagentnode** parameter in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file (`tdpexc.cfg`) or (`tdpsql.cfg`) by default). To specify this parameter with the **Properties** window of Microsoft Management Console (MMC), select VSS backup. In the **Properties** window, you can update the Local DSMAGENT Node name. Otherwise, use the **tdpexc set** or **tdpsqlc set** command to specify this parameter.

Agent node - Remote DSMAGENT Node

This node name is a separate system that must also have the backup-archive client, and the VSS provider installed. In addition, for Exchange Server workloads, ensure that you install the same level of the Exchange System Management Tools that is installed on your Exchange production server. This node is responsible for moving VSS snapshot data from local shadow volumes to the IBM Spectrum Protect server.

This node name is referred to as the Remote DSMAGENT Node and is specified with the **remotedsmagentnode** parameter in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file (`tdpexc.cfg` or `tdpsql.cfg` by default). To specify this parameter with the **Properties** window of MMC, select VSS backup. Then, you can update the Remote DSMAGENT Node name. Otherwise, use the **tdpexc set** or **tdpsqlc set** command to specify this parameter.

The choice of available systems depends on whether the systems have access to the local shadow volumes that contain the VSS snapshot backups. This node name is only valid for VSS environments that support shadow copies that can be transported.

If you are using the default VSS system provider, you cannot specify the node name.

If you are running offloaded backups, a dedicated (and unique) remote DSM agent node must exist for each local DSM agent node.

Ensure that the **localdsmagentnode** and **remotedsmagentnode** are registered to the same IBM Spectrum Protect server that is specified in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server options file (`dsm.opt`) and the backup-archive client options file (also `dsm.opt`).

Configuring the system that runs offloaded backups

Complete the following steps on the computer that is running the offloaded backups. This task is for VSS operations only.

Procedure

1. Configure the IBM Spectrum Protect backup-archive client if it is not already configured.
If the backup-archive client is already configured, you can use existing client services. Select **Utilities > Setup Wizard > Help me configure the TSM Backup Archive Client**.
The node name for this system is called the Remote DSMAGENT Node and is specified with the **remotedsmagentnode** parameter in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file on the local system.
2. Install and configure the IBM Spectrum Protect Client Acceptor (CAD) Service and the Remote Client Agent Service (DSMAGENT) if these services are not already installed.
If a client CAD Service is already installed and configured, you can use an existing one. Use the backup-archive client Setup wizard to guide you through the CAD installation process by selecting **Utilities > Setup Wizard > Help me configure the TSM Web Client**.
3. Install the Microsoft Exchange Server management tools from the Microsoft Exchange Server installation media. Take note of the Microsoft Exchange Server Management tools binary directory, for example, `C:\Program files\Exchsrvr\bin`. Verify that the `ESEUTIL.EXE` tool is stored in this directory. IBM Spectrum Protect Snapshot for Microsoft Exchange Server uses this tool to run automatic integrity checking of the VSS backup.
Also, the Exchange Server does not need to be installed or running on this system. Only the Microsoft Exchange Server management tools must be installed on this system. For more information about the necessary license requirements, see the Microsoft Exchange Server documentation.
4. Add the Microsoft Exchange Server binary path to the **PATH** statement in the system environment variables.
For example:

```
"C:\Program files\Exchsrvr\bin"
```
5. Install and configure a VSS provider if you do not use the default system VSS provider. Consult the VSS provider documentation for information about the configuration of that software.

Configuring to protect SQL Server data

With IBM Spectrum Protect Snapshot for SQL Server, you can configure SQL Server clustered environments, high availability environments, and Windows Server Core environments.

Configuring IBM Spectrum Protect Snapshot for SQL Server clustered environments

Depending on the applications that are installed, you can configure IBM Spectrum Protect Snapshot to operate in clustered environments with SQL Server, or file system workloads.

Related tasks

[Troubleshooting configuration errors in a failover clustered environment](#)

If you encounter errors when you configure a failover clustered environment, determine whether the options in the backup-archive client and application-specific `dsm.opt` files are specified correctly.

Configuring IBM Spectrum Protect Snapshot for SQL Server stand-alone configuration in a SQL failover cluster environment with shared disks or cluster shared volumes

You can configure IBM Spectrum Protect Snapshot as a stand-alone configuration, and protect SQL Server workloads in a clustered environment that uses either shared disks or cluster shared volumes (CSV).

About this task

In the backup-archive client `dsm.opt` file, each system uses its node name as the local agent node for IBM Spectrum Protect Snapshot. The same IBM Spectrum Protect Snapshot for SQL Server node name, `VirtualClusterNodeName`, is applied to all of the systems in the cluster.

The `VSSALTSTAGINGDIR` path must point to an accessible directory on a shared disk that all cluster nodes can access, for example, `X:\vss_staging`. The `VSSALTSTAGINGDIR` option must be specified in the backup-archive client options file, `baclient\dsm.opt`, and in the IBM Spectrum Protect Snapshot for SQL Server options file, `tdpsql\dsm.opt`, and the option argument must be the same. For example, `X:\vss_staging` is the absolute path to the VSS staging directory on a shared disk that all cluster nodes can access.

You also specify the `clustersharedfolder` option in the IBM Spectrum Protect Snapshot for SQL Server options file. This option specifies the directory location in which to store encrypted password files and the path must point to a directory location that is shared by all nodes in the cluster. If a failover occurs, the backup-archive client uses this option to determine where the password files are located.

Note:

You can set the `clustersharedfolder` option value to be the same as that set for the `VSSALTSTAGINGDIR` directory path.

Alternatively, you can go to the `baclient` directory and open either the `dsm.opt` or `custom.opt` file when the configuration wizard mode is in either default or custom mode. By default, the VSS Requestor `dsm.opt` file is in the IBM Spectrum Protect backup-archive client installation directory, `c:\Program Files\tivoli\tsm\baclient\dsm.opt`. The location of the IBM Spectrum Protect Snapshot for SQL Server options file is `c:\Program Files\tivoli\tsm\TDPSQL\dsm.opt`.

Procedure

1. Install

IBM Spectrum Protect Snapshot for SQL on all cluster nodes, and in the same directory on all cluster nodes. The default installation directory is `c:\Program Files\tivoli\FlashCopyManager\`.

2. On the active node in the cluster, run the stand-alone configuration wizard, from the Microsoft Management Console (MMC), and follow these steps:

- a) Specify the same IBM Spectrum Protect node name for the IBM Spectrum Protect Snapshot for SQL Server configuration.
- b) Specify the same folder location for the VSS alternative staging directory to point to an existing directory on a shared disk, or on a cluster shared volume, for example, `X:\vss_staging`.

You can use the remote management configuration feature for remote management of other cluster nodes.

3. After the configuration wizard finishes successfully, close MMC.
4. Configure the client acceptor daemon.

For more information, see [Configuring cluster protection \(Windows Server clients\) \(http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.7/client/t_cfg_clus_wizard_win2008.html\)](http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.7/client/t_cfg_clus_wizard_win2008.html)

5. Repeat these steps on all other nodes in the cluster.

Important: Each node must be the active node in the cluster when the stand-alone configuration wizard is running.

6. Verify that the configuration is valid by manually checking every cluster node as follows:
 - a) Open MMC and select **Dashboard > Configuration > Files**.
 - b) Depending on whether the configuration wizard mode is in either default or custom mode, browse to either the `dsm.opt` or `custom.opt` file. Ensure that the settings in the file are as follows:

```
NODename localdsmagent
PASSWORDAccess generate
TCPServeraddress flashcopymanager
CLUSTERnode no
CLUSTERDISKSONly no
VSSALTSTAGINGDIR X:\vss_staging
```

- c) Select `dsm.opt` under the IBM Spectrum Protect Snapshot for SQL Server section, and ensure that the settings are as follows:

```
NODename VirtualClusterNodeName
PASSWORDAccess generate
TCPServeraddress flashcopymanager
CLUSTERnode yes
VSSALTSTAGINGDIR X:\vss_staging
CLUSTERSHAREDFOLDER X:\vss_staging
```

Related tasks

[“Troubleshooting configuration errors in a failover clustered environment” on page 181](#)

If you encounter errors when you configure a failover clustered environment, determine whether the options in the backup-archive client and application-specific `dsm.opt` files are specified correctly.

Configuring IBM Spectrum Protect Snapshot for file system and custom applications in a Microsoft Cluster Server environment

You can configure IBM Spectrum Protect Snapshot to support a file system and custom application workload in a Microsoft Cluster Server environment.

Before you begin

- Before you begin your configuration, read the entire procedure.
- Perform the configuration steps in the same way on all of the nodes in the cluster.
- Keep the number of cluster groups to a minimum. If possible, have only one cluster resource group that contains all physical disk resources.
- Use a dedicated volume (VSS staging directory volume) for each cluster resource group. The VSS staging directory volume must have a minimum of 1 GB of storage space for IBM Spectrum Protect Snapshot metadata. This volume must be part of the cluster resource group and must be able to fail over with the cluster resource group. This VSS staging directory volume must not be one of the volumes that is backed up with IBM Spectrum Protect Snapshot.
- Use the `clustersharedfolder` option in the IBM Spectrum Protect Snapshot for SQL Server options file (`tdpsql\dsm.opt`) to specify the directory location in which to store the encrypted password file. This ensures that any update to the password file is on a resource that is shared among all nodes in the cluster. For more information, see [Clustersharedfolder](#)

About this task

Complete the following steps for each node on the cluster.

Procedure

1. Install IBM Spectrum Protect Snapshot.
2. Start Microsoft Management Console (MMC).

You are automatically prompted to run the configuration wizard. If the configuration wizard prompt is not displayed automatically, expand the following tree nodes: **IBM Spectrum Protect > Dashboard > Manage > Configuration**. Then, select **Wizards**. Depending on your environment, launch either the **Standalone Configuration** or IBM Spectrum Protect **Configuration** wizard. If you have an IBM Spectrum Protect server, select IBM Spectrum Protect **Configuration**, otherwise, if you do not have access to an IBM Spectrum Protect server, select **Standalone Configuration**.

3. Select the **File System** checkbox. To start the wizard, click **Next**.

Table 14. Configuration options for file system and custom applications	
Configuration	Action
Standalone Configuration	<p>a. On the requirements check pane, you might see a Warnings link next to the Cluster check rule. Click Warnings and MMC displays the Issue Resolution Windows for vssaltstagingdir path. In the Path field, enter the path of your VSS staging directory volume. If you are configuring multiple resource groups, start with the dedicated volume that belongs to the cluster resource group that you want to configure first.</p> <p>b. Complete the wizard pages.</p> <p>c. Click Finish to complete the initial configuration.</p> <p>d. Exit MMC.</p> <p>After the configuration wizard completes, the following are the contents of the different options files. Values might differ slightly:</p> <ul style="list-style-type: none"> • In the backup-archive client options file: <pre>NODename OHELLO PASSWORDAccess generate TCPServeraddress flashcopymanager HTTPport 1581 CLUSTERnode no CLUSTERDISKSONly no VSSALTSTAGINGDIR J:\vssaltstagingdir</pre> • In the IBM Spectrum Protect Snapshot file system options file: <pre>NODename OHELLO_FS PASSWORDAccess generate TCPServeraddress flashcopymanager HTTPport 1581 CLUSTERnode yes VSSALTSTAGINGDIR J:\vssaltstagingdir CLUSTERSHAREDFOLDER J:\vssaltstagingdir</pre>

Table 14. Configuration options for file system and custom applications (continued)

Configuration	Action
IBM Spectrum Protect Configuration	<p>a. On the requirements check pane, you might see a Warnings link next to the Cluster check rule. Click Warnings and MMC displays the Issue Resolution Windows for vssaltstagingdir path. In this Path field, enter the path of your VSS staging directory volume. If you are configuring multiple resource groups, start with the dedicated volume that belongs to the cluster resource group you want to configure first.</p> <p>b. Complete the wizard pages.</p> <p>c. Click Finish to complete the initial configuration. After the configuration wizard is complete, the following contents are displayed in the different options files. Values might differ slightly. In the following OPT files, ensure that the DSM agent node and IBM Spectrum Protect Snapshot node are registered and granted proxy.</p> <ul style="list-style-type: none"> In the backup-archive client options file: <pre data-bbox="618 558 1187 737"> NODename OTHELLO PASSWORDAccess generate TCPServeraddress orion.storage.usca.ibm.com TCPPort 1500 CLUSTERnode no CLUSTERDISKSONly no VSSALTSTAGINGDIR J:\vssaltstagingdir </pre> In the IBM Spectrum Protect Snapshot file system options file: <pre data-bbox="618 800 1187 989"> NODename CLUSTER_FS PASSWORDAccess generate TCPServeraddress orion.storage.usca.ibm.com TCPPort 1500 HTTPport 1581 CLUSTERnode yes VSSALTSTAGINGDIR J:\vssaltstagingdir CLUSTERSHAREDFOLDER J:\vssaltstagingdir </pre> <p>d. Exit MMC.</p> <p>e. Open a Windows command line and change the directory to the backup-archive client directory location. Default location: c:\Program Files\Tivoli\tsm\baclient</p> <p>f. To connect to the IBM Spectrum Protect server, enter the dsmc command. You might need to provide your user ID and password for the backup-archive client DSMAGENT node to save the password on the registry. Exit the dsmc.</p> <p>g. Using the same Windows command line, enter the following command to connect to the IBM Spectrum Protect server by using the IBM Spectrum Protect Snapshot node:</p> <pre data-bbox="602 1272 1325 1325"> dsmc -optfile="c:\Program Files\Tivoli\FlashCopyManager\dsm.opt" </pre> <p>You might need to provide your user ID and password for the IBM Spectrum Protect Snapshot node to save the password on the registry. Exit the dsmc and exit the Windows command line.</p>

4. (Standalone Configuration only) Complete the following steps:

- Open the Windows Services MMC. Stop both the CAD and Agent Services that are named, by default, *TSM Client Acceptor* and *TSM Remote Client Agent*.
- Open a Windows command line and change directories to the IBM Spectrum Protect Snapshot installation directory. The default location: C:\Program Files\Tivoli\FlashCopyManager
- Open the `dsm.opt` file by using Notepad, and change the `nodename` option to a different name that would best describe your cluster.
For example: `NODename cluster_fs`
- Specify the VSS staging directory volume for IBM Spectrum Protect Snapshot VSS metadata.
Add the **vssaltstagingdir** *path* option at the end of the file. The path must be the path of your VSS staging directory volume. If you are configuring multiple resource groups, start with the dedicated volume in the cluster resource group that you want to configure first.

For example, if the J: drive is the dedicated VSS staging directory volume in the cluster resource that you want to configure, specify this option:

```
VSSALTSTAGINGDIR J:\vssaltstagingdir
```

- e) Save and close the IBM Spectrum Protect Snapshot options file.
- f) Change the directory to the backup-archive client installation directory. Default location:
C:\Program Files\Tivoli\tsm\baclient
- g) Open the `dsm.opt` file by using Notepad, and add the exact same value for the **vssaltstagingdir** option as exists in the IBM Spectrum Protect Snapshot `dsm.opt` file.
For example:

```
VSSALTSTAGINGDIR J:\vssaltstagingdir
```

- h) Specify the `clustersharedfolder` option to specify the directory location in which to store encrypted password files. Specify a location that all nodes in the cluster nodes can access. If a failover occurs, the backup-archive client uses this option to determine where the password files are located. You can set the `clustersharedfolder` option value to be the same as that set for the `VSSALTSTAGINGDIR` directory path.
 - i) Save and close the backup-archive client options file.
 - j) Exit the Windows command line.
5. If this node is the first cluster node that you are configuring, open the Microsoft Failover Cluster Manager.
- Go to the cluster resource group that you are configuring. Right-click the resource group and select **Add a resource > Generic Service**. From the **New Resource Wizard** dialog that is displayed, select the CAD service that is named, by default, *TSM Client Acceptor*, and complete the wizard configuration. From the Failover Cluster Manager, bring the *TSM Client Acceptor* service resource online. If you are repeating this procedure for another cluster node, the *TSM Client Acceptor* service resource is already configured. In this scenario, go to the resource group and open the service resource online.
6. Open a Windows command line and change directories to the IBM Spectrum Protect Snapshot installation directory.

Default location: C:\Program Files\Tivoli\FlashCopyManager

Enter the following command:

```
fccli query component
```

A list of all available volumes for backup is displayed.

Your first group resource is configured.

7. Complete the same procedure on the other nodes in your cluster.

Before you begin the procedure on other nodes in the cluster, go to the Windows Services MMC, and stop the *TSM Remote Client Agent* service if it is running. Then, using the Microsoft Failover Cluster Manager, make the *TSM Client Acceptor* service offline from the resource group that you configured. Finally move the group resource that you configured to the other node in the cluster and restart the procedure. A final note: All nodes in the cluster must have identical IBM Spectrum Protect Snapshot options file. For the backup-archive client `dsm.opt` file, each node in the cluster can use its own node name, but everything else in the options file must be identical.

Related concepts

[Prerequisites](#)

To create VSS snapshot backups of NTFS or ReFS file systems and applications, use IBM Spectrum Protect Snapshot. When you back up applications and file systems, IBM Spectrum Protect Snapshot must access the data.

Moving standard SQL databases to the AlwaysOn node

You can specify the `/USEALWAYSONnode` parameter with the `backup` command to back up standard SQL databases to the file space for the AlwaysOn node. This transition can make it easier for you to manage all your database backups under a single node name.

About this task

If you want to regularly back up standard SQL databases to the file space for the AlwaysOn node, you can use the `set` command.

The AlwaysOn node name is required when you configure IBM Spectrum Protect Snapshot with SQL Server 2012 and later versions. It is not necessary to specify the AlwaysOn node name during each backup, query, or restore operation of an availability database.

The AlwaysOn node does not affect where standard databases are backed up. The standard databases continue to be backed up to the IBM Spectrum Protect Snapshot node unless the `/USEALWAYSONnode` option is specified.

Procedure

- Enter the following command to back up your standard SQL databases to the file space for the AlwaysOn node:

For example,

```
TDPSQLC Backup *|dbname[,dbname,...] Full /USEALWAYSONnode
```

You can use the wildcard character (*) to back up all databases, or specify a list of database names that are separated by commas.

For example:

```
TDPSQLC Backup standard_db01,standard_db02 Full /USEALWAYSONnode
```

Configuring availability replicas to run scheduled data backups

When an availability database is replicated across multiple availability replicas in an availability group, a configuration option is available that you can use to select a single replica on which to run a backup operation instead of backing up all replicas.

About this task

Microsoft SQL Server 2012 and later versions provide a set of configuration options that you can use to specify whether scheduled backups are run on the primary or secondary availability replica. You can use the IBM Spectrum Protect Snapshot GUI to set these options.

The configuration option can also be used to offload the backup from a primary replica to a secondary replica for load balancing. When databases fail over, backups must continue to run from other replicas to ensure that high availability is maintained.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the **Management** section of the window, click **Protect Data** next to the SQL Server workload.
3. In the **Action** pane, click **Properties**.
4. Click the **AlwaysOn Preferences** property page.
5. In the **Availability group** field, select the **AlwaysOn Availability Group** for which you want to set up backup preferences.

6. In the **Preferred replica** field, select which replica is the preferred replica on which to run scheduled backups.
 - Select **Prefer Secondary replica** if you want scheduled backups to occur on a secondary replica, if it is available. Otherwise, use the primary replica for the scheduled backup.
 - Select **Secondary only** if you want scheduled backups to occur only on a secondary replica.
 - Select **Primary** if you want scheduled backups to occur only on the primary replica.
 - Select **Any replica** if you want scheduled backups to occur on any availability replica.
7. For each availability replica that is listed in the Availability replicas list box, specify whether it is a candidate for running scheduled backups by specifying the backup priority for that replica. A value of 1 has the lowest priority, and a value of 100 has the highest priority. A value of 0 indicates that the replica is excluded from schedule backup operations.
8. Click **OK** to save your configuration and exit the **Data Protection Properties** page. The settings are saved to the `tdpsql.cfg` file and can be replicated to the other replicas in the availability group.

What to do next

After you configure where scheduled backups are run, the administrator can specify the **tdpsql backup** command along with the **/ALWAYSONPriority** parameter in a backup schedule. For example:

```
tdpsqlc backup TestDb1 full /ALWAYSONPriority
```

When this scheduled backup command is run, IBM Spectrum Protect Snapshot queries the SQL Server to determine the highest-priority availability replica that is active or online, ordered by preference. If the replica meets the specified criteria, the replica is backed up. Otherwise, the backup operation ends and a message is added to the log to indicate why the replica was not backed up.

An administrator can create a common backup schedule to run on all availability replicas. When the backup schedule starts, each **tdpsqlc** command queries each replica to determine whether it is to run the backup. Only one of the scheduled backups runs the backup.

Configuring IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core

You can manually configure IBM Spectrum Protect Snapshot to protect your SQL Server 2012 and later versions on Windows Server Core.

Before you begin

Add the remote server core system to Microsoft Management Console (MMC) assuming that IBM Spectrum Protect Snapshot is already installed on the remote system and you configured PowerShell remoting on the remote system, on either a local or centralized computer. On the local system that runs MMC, point to the server core system and complete the **TSM Configuration** wizard. By using the wizard, you can provision and configure the software as if the wizard is running on the remote system.

If you do not use the remote configuration option, ensure that you install IBM Spectrum Protect Snapshot and the IBM Spectrum Protect backup-archive client on the system that runs the Microsoft SQL Server.

Procedure

1. Create a node on the IBM Spectrum Protect server for the backup-archive client and IBM Spectrum Protect Snapshot.
 - If you are protecting availability databases in an AlwaysOn Availability Group, you must also create the AlwaysOn node on the IBM Spectrum Protect server.
2. If you intend to run offloaded VSS backups, set up a remote node to run the offloaded backup operation on a remote computer.
3. Configure the backup-archive client options file (`dsm.opt`).
4. Configure the IBM Spectrum Protect Snapshot for Microsoft SQL Server option files (`dsm.opt` and `tsmfcm.cfg`).

5. If you use IBM Spectrum Protect policy sets, specify a management class to use for your IBM Spectrum Protect Snapshot backups.

Creating a node on the IBM Spectrum Protect server

After you install the IBM Spectrum Protect client and IBM Spectrum Protect Snapshot, you must set up a node name and password and register your node with the IBM Spectrum Protect server. When a new node is registered, an administrative user ID must be created for the node. The IBM Spectrum Protect server administrator must specify the `userid` option with the **register node** command.

About this task

When you register your node, you create a file space on the IBM Spectrum Protect server where the backups of your data are stored. You must set up a client node and an IBM Spectrum Protect Snapshot node. If you are protecting availability databases in an AlwaysOn Availability Group, you must also register the AlwaysOn node.

Follow these procedures if you installed the IBM Spectrum Protect administrative command line client. If you did not install the administrative client, the nodes must be registered on the IBM Spectrum Protect server.

Procedure

1. Start an administrative client session by entering the following command at the command line:

```
C:\Program Files\Tivoli\TSM\baclient\dsmadm
```

2. To register a client node, enter the following command:

```
reg node client_nodename password backdel=yes userid=client_nodename
```

where *client_nodename* is the node name for the client and *password* is the password that you want to use for the client. The **backdel=yes** parameter indicates that you can delete backup objects in your file space on the server. The node name and administrative user ID must be the same.

For example:

```
reg node doomvm3 doomvm3passwd backdel=yes userid=doomvm3
```

3. To register an IBM Spectrum Protect Snapshot for Microsoft SQL Server node, enter the following command:

```
reg node sql_nodename password backdel=yes userid=sql_nodename
```

where *sql_nodename* is the node name for the IBM Spectrum Protect Snapshot for Microsoft SQL ServerData Protection node and *password* is the password to use for the SQL node. The **backdel=yes** parameter indicates that you can delete backup objects in your file space on the server. The node name and administrative user ID must be the same.

For example:

```
reg node doomvm3_sql doomvm3sqlpasswd backdel=yes userid=doomvm3_sql
```

Tip: To easily identify the node as a node for IBM Spectrum Protect Snapshot for Microsoft SQL Server, add "_sql" to the end of the node name.

4. To register the AlwaysOn node, enter the following command:

```
reg node alwayson_nodename password backdel=yes userid=alwayson_nodename
```

where *alwayson_nodename* is the name for the AlwaysOn node and *password* is the password that you want to use for the AlwaysOn node. The **backdel=yes** parameter indicates that you can delete backup objects in your file space on the server. The node name and administrative user ID must be the same.

For example:

```
reg node myalwaysonnode alwaysonpasswd backdel=yes userid=myalwaysonnode
```

What to do next

To use IBM Spectrum Protect server policy sets, the IBM Spectrum Protect must define the policy domains, policy sets, management classes, copy groups, and storage pools.

These definitions are necessary to meet your IBM Spectrum Protect Snapshot backup and restore requirements. For VSS operations, IBM Spectrum Protect server authentication must be on.

Setting up a proxy node for offloaded VSS backups in the Windows Server Core environment

If you want to offload VSS backups to the IBM Spectrum Protect Snapshot, you must define a remote node to run the offloaded backups. This step is part of the configuration tasks for operating IBM Spectrum Protect Snapshot on Windows Server Core.

About this task

IBM Spectrum Protect Snapshot can offload VSS backups by using a remote computer to create the backup instead of using the local computer. To run an offload backup by using a remote node, you must first set the remote node as an agent of the local IBM Spectrum Protect Snapshot node.

If you are protecting availability databases in an AlwaysOn Availability Group, you must set the remote node as an agent of the AlwaysOn node.

Before you begin, ensure that the IBM Spectrum Protect client is installed and configured on the remote computer.

Procedure

To define the proxy node relationship, the IBM Spectrum Protect administrator can enter the **grant proxynode** command from the IBM Spectrum Protect server administrative console.

- For standard IBM Spectrum Protect Snapshot nodes, enter the following command:

```
grant proxynode target=local_sql_node agent=remote_node
```

where *local_sql_node* is the node name of the local IBM Spectrum Protect Snapshot node, and *remote_node* is the remote IBM Spectrum Protect client node that runs the remote backups.

For example:

```
grant proxynode target=doomvm3_sql agent=babar
```

- For AlwaysOn nodes, enter the following command:

```
grant proxynode target=alwayson_node agent=remote_node
```

where *alwayson_node* is the name of the AlwaysOn node, and *remote_node* is the remote IBM Spectrum Protect client node that runs the remote backups.

For example:

```
grant proxynode target=myalwaysonnode agent=babar
```

- To display the client nodes with authority to act as proxy to other clients, run the following command from the administrative console of the server:

```
query proxynode
```

Configuring the client in the Windows Server Core environment

You must configure the IBM Spectrum Protect client node that you created. This step is part of the initial configuration tasks before you can use IBM Spectrum Protect Snapshot in the Windows Server Core environment.

About this task

You must configure the client options file (`dsm.opt`), set the environment variables, and install and setup the IBM Spectrum Protect client acceptor service and remote client agent service.

Procedure

1. Configure the client options file:

a) Change to the backup-archive client installation directory.

For example, issue the following command in a **Command Prompt** window:

```
cd C:\Program Files\Tivoli\TSM\baclient
```

b) Open the `dsm.opt` file with a text editor and enter the following statements:

```
PASSWORDACCESS GENERATE
COMMMethod TCPip
TCPport 1500
nodename client_nodename
TCPSEVERADDRESS tsm_server
```

The following list contains brief explanations of the client options in the statements:

PASSWORDACCESS GENERATE

Instructs the client to save the password whenever the `/tsmpassword` option is used so that you do not have to enter the password with every command.

TCPport 1500

Specifies that the client accesses the IBM Spectrum Protect server at TCP/IP port 1500. 1500 is the default port number.

nodename client_nodename

Specifies the newly created node for the backup-archive client.

TCPSEVERADDRESS tsm_server

Specifies the name of the IBM Spectrum Protect server. You can enter the server IP address or the fully qualified domain name.

For example:

```
NODename DOOMVM3
PASSWORDAccess generate
TCPServeraddress gijoe
TCPport 1500
```

2. Install and start the IBM Spectrum Protect client acceptor service and remote client agent service.

a) Install the client acceptor service by entering the following command in a **Command Prompt** window:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install cad
/name:"servicename" /node:nodename /password:password
/autostart:yes
```

where `nodename` is the client node name, `password` is the client password, and `servicename` is the name that you want to use for the client acceptor service. The default name is "TSM Client Acceptor".

For example:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install cad /name:"TSM CAD"
/node:DOOMVM3 /password:doomvm3passwd /autostart:yes
```

- b) Install the remote client agent service by entering the following command in a **Command Prompt** window:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install remoteagent
/name:"servicename" /node:nodename /password:password
/partnername:"partner service name"
```

The node name for the IBM Spectrum Protect Client Acceptor and the Remote Client Agent must be set to the VSS requestor node. The default service name is "TSM Remote Client Agent". The value for the **/partnername** option must match the name of the client acceptor service that you created. The default name is "TSM Client Acceptor".

For example:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install remoteagent
/name:"TSM AGENT" /node:DOOMVM3 /password:doomvm3passwd
/partnername:"TSM CAD"
```

- c) Start the client acceptor service by entering the following command:

```
net start "servicename"
```

where *servicename* is the name of the client acceptor service that you created. For example:

```
net start "TSM CAD"
```

Do not start the remote client agent service manually. The remote client agent service is automatically started by the client acceptor service when it is needed.

Configuring IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core

You must configure IBM Spectrum Protect Snapshot before you can protect your Microsoft SQL Server 2012 and later databases in the Windows Server Core environment.

Before you begin

Restriction: You cannot specify the following special characters in the SQL Server database name on IBM Spectrum Protect Snapshot:

- Question character (?)
 - Multibyte character (,)
 - Multibyte character (^)
 - Asterisk character (*)
 - Colon character (:)
 - Backslash character (\)
- cannot be used in IBM Spectrum Protect Snapshot version 4.1.0 or earlier versions
- cannot be used in IBM Spectrum Protect Snapshot version 4.1.0 or earlier versions

About this task

You must configure the client options file (*dsm.opt*) and IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file (*tdpsql.cfg*).

Procedure

1. Edit the client options file (*dsm.opt*).

- a) In the IBM Spectrum Protect Snapshot installation directory, open the client options file (*dsm.opt*) with a text editor.
- b) Add the following statements to the client options file:

```
NODename          sql_nodename
PASSWORDAccess    Generate
COMMMethod        TCPip
TCPServeraddress  tsm_server
```



```
TCPPort      1500
TCPWindowSize 63
TCPBuffSize  32
```

where **nodename** is the IBM Spectrum Protect Snapshot node name, and **TCPServeraddress** is the name of the IBM Spectrum Protect server. You can enter the server IP address or the fully qualified domain name.

For example:

```
NODename DOOMVM3_SQL
PASSWORDAccess generate
TCPServeraddress gijoe
TCPPort 1500
```

2. Edit the `tdpsql.cfg` file.

- a) In the IBM Spectrum Protect Snapshot installation directory, open the configuration file (`tdpsql.cfg`) with a text editor.
- b) Add the following statements in the `tdpsql.cfg` file:

```
SQLSERVer      sql_server
FROMSQLserver  sql_server
SQLAUTHentication  INTEGRated
MOUNTWaitfordata  Yes
BACKUPMethod     Legacy|VSS]
DIFFESTimate     20
BUFFers          3
BUFFERSize       1024
STRIPes          1
SQLBUFFers       0
SQLBUFFERSize    1024
LOGPrune         60
LANGuage         ENU
BACKUPDestination [LOCAL|TSM|BOTH]
LOCALDSMAgentnode local_node
REMOTEDSMAgentnode remote_node
ALWAYSONNode     alwayson_node
USEALWAYSONnode [Yes|No]
ENABLEREplacementchars [Yes|No]
LOGFile          tdpsql.log
```

The options in the `tdpsql.cfg` file are as follows:

SQLSERVer

Specifies the name of the Microsoft SQL Server that is running on the local computer.

BACKUPMethod

Determines whether to run a legacy or VSS backup.

BACKUPDestination

Determines whether to run a local backup, IBM Spectrum Protect backup, or both. For legacy backups, only IBM Spectrum Protect is used.

LOCALDSMAgentnode

Specifies the local node name of the client that is running on the local computer. This option is required for VSS offloaded backups.

REMOTEDSMAgentnode

Specifies the remote client node that runs the VSS offloaded backups on a remote computer.

ALWAYSONNode

Specifies the IBM Spectrum Protect node name that is used to back up availability databases in an AlwaysOn Availability Group.

USEALWAYSONnode

Specify *Yes* to set the AlwaysOn node as the default node for all backup operations of standard and availability databases. You can use this option to change database backups from a standard IBM Spectrum Protect Snapshot node to an AlwaysOn node.

Specify *No* to back up standard databases to the IBM Spectrum Protect Snapshot node. Availability databases are always backed up with the AlwaysOn node.

ENABLEREPlacementchars

Specify *Yes* to enable IBM Spectrum Protect Snapshot to process backslash (\) or colon (:) characters in a database name, and back up the database to IBM Spectrum Protect.

Specify *No* to prevent database backups to IBM Spectrum Protect if a user-defined string is substituted for a backslash (\) or colon (:) character in the database name.

Restriction: The **ENABLEREPlacementchars** parameter applies only to IBM Spectrum Protect Snapshot version 4.1.1 and later versions. The maximum length of the database name is 128 characters.

3. If you run the stand-alone configuration on IBM Spectrum Protect Snapshot, complete the following steps:
 - a) In the IBM Spectrum Protect Snapshot installation directory, open the client options file (`dsm.opt`) with a text editor.
 - b) Edit the `dsm.opt` file and change the `TCPServeraddress tsm_server` statement to `TCPServeraddress flashcopymanager`.
 - c) If installed, remove the IBM Spectrum Protect client acceptor service. Run the following command from a **Command Prompt** window:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil remove /name:"TSM CAD"
```

In this case, TSM CAD is the name of the client acceptor service that you want to remove.

- d) Reinstall the remote client agent service by entering the following command:

```
C:\Program Files\Tivoli\TSM\baclient\dsmcutil install remoteagent  
/name:"TSM AGENT" /node:DOOMVM3_SQL /password:doomvm3sqlpasswd  
/partnername:"TSM CAD"
```

In this case, TSM AGENT is the name of the remote agent and TSM CAD is the name of the client acceptor service.

- e) Start the client acceptor service by entering the following command:

```
net start "TSM CAD"
```

In this case, TSM CAD is the name of the client acceptor service are starting.

Do not start the remote client agent service manually. The remote client agent service is automatically started by the client acceptor service when it is needed.

4. Optional: Use the **VSSPOLICY** option to specify a management class for VSS backups.

Unless specified otherwise, IBM Spectrum Protect Snapshot uses the default management class of the policy domain that its node name is in. To specify that IBM Spectrum Protect Snapshot uses a different management class, add the **VSSPOLICY** option to the `tdpsqlc.cfg` file. The format of the option is as follows:

```
VSSPOLICY SQL_server_name "db_name" backup_type backup_dest mgmt_class
```

For example:

```
VSSPOLICY doomvm3 * FULL LOCAL MGMT2
```

This statement specifies that IBM Spectrum Protect Snapshot uses the management class MGMT2 for local backups of any database in the SQL Server named `doomvm3`.

Configuring your system for mailbox restore operations

To use IBM Spectrum Protect Snapshot for Exchange Server to restore mailboxes and mailbox items, you must complete the configuration wizard.

About this task

The IBM Spectrum Protect Snapshot for Exchange Server configuration wizard verifies that user permissions and software versions are correct.

- Ensure that you have the role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- **Exchange Server 2013:** Install the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects on the Exchange server from which you are running the mailbox restore operations.

Tip: Do not install Microsoft Outlook 2010 or 2013 on the same server that IBM Spectrum Protect Snapshot for Exchange Server uses for mailbox restore operations. Conflicts might occur in the MAPI configurations.

- **Exchange Server 2016 or later:** The mailbox restore operation in Mailbox Restore Browser view uses Microsoft 32-bit Outlook 2016 or later versions as the MAPI client. Microsoft does not support installations of Outlook on the same machine as Exchange server. It is recommended that Outlook is installed on a separate machine. IBM Spectrum Protect Snapshot must be installed on both the Outlook machine and the Exchange Server machine. With IBM Spectrum Protect Snapshot installed on the Outlook machine, you can open the **Mailbox Restore Browser** view of the remote Exchange server using Remote Management. The mailbox restore operation in **Mailbox Restore** view does not require Outlook or other MAPI client. These operations can be performed in IBM Spectrum Protect Snapshot on the Exchange server directly.

Tip: Ensure that the logon user's mailbox is in a database on Exchange Server 2016 or later version.

Procedure

1. If you are using an incorrect Microsoft MAPI Client version, click the **Warnings** link and install the correct version.
2. If you do not have all the management roles for individual mailbox restore operations, click the **Warnings** link and follow the wizard prompts to add the missing Exchange roles.
If you are a member of the Exchange Organization Management group, you can automatically add the missing roles. If you are not a member of the Exchange Organization Management group, you must manually add the missing roles.
3. For Exchange Server 2013, configure the Client Access Server (CAS) role to run Mailbox Restore operations. For more information about specifying the CAS with the **set** command, see the **Set syntax** command.

Related concepts

[Exchange mailbox restore operations](#)

By using IBM Spectrum Protect Snapshot, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

Configuring your system for mailbox restore operations (Exchange 2016 and later)

To use the IBM Spectrum Protect Snapshot for Exchange Server to restore mailboxes and mailbox items with Exchange Server 2016 and later versions, you must configure the system for mailbox restore operations.

About this task

The IBM Spectrum Protect Snapshot for Exchange Server configuration wizard verifies that user permissions and software versions are correct.

- Ensure that you have the role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- The mailbox restore operation in **Mailbox Restore Browser** view uses Microsoft 32-bit Outlook 2016 or later versions as the MAPI client. Microsoft does not support installations of Outlook on the same machine as Exchange server. It is recommended that Outlook is installed on a separate machine.
- IBM Spectrum Protect Snapshot must be installed on both the Outlook machine and the Exchange Server machine.
- With IBM Spectrum Protect Snapshot installed on the Outlook machine, you can open the **Mailbox Restore Browser** view of the remote Exchange server using Remote Management.
- The mailbox restore operation in **Mailbox Restore** view does not require Outlook or other MAPI client. These operations can be performed in IBM Spectrum Protect Snapshot on the Exchange server directly.

Tip: Ensure that the logon user's mailbox is in a database on Exchange Server 2016 or later version.

Procedure

1. Install Microsoft 32-bit Outlook 2016 or later version on a separate machine without Exchange server. The Outlook machine should be in the same domain as Exchange Server.
2. Install IBM Spectrum Protect Snapshot on both the Outlook machine and the Exchange server.
3. Use the **Mailbox Restore Only** wizard to configure IBM Spectrum Protect Snapshot on the Outlook machine. Make sure all requirements are met in the wizard.
4. Configure remote management to use IBM Spectrum Protect Snapshot on the Outlook machine to manage the Exchange server remotely. For more information, see [“Enabling Windows PowerShell Remoting for Remote Management and Remote Mounting”](#) on page 155.
5. In IBM Spectrum Protect Snapshot on the Outlook machine, open the **Mailbox Restore Browser** view of the remote Exchange Server node to perform mailbox restore operations.

Examples of distributed VSS backups in Microsoft Exchange Database Availability Groups and Microsoft SQL AlwaysOn Availability Groups

To use IBM Spectrum Protect Snapshot for VSS backups, see the following illustration of a sample deployment. The illustration applies to Microsoft Exchange Database Availability Groups (DAGs) and Microsoft SQL AlwaysOn Availability Groups (AAGs).

In the following illustration, an Exchange Server DAG is shown, but you can use an SQL Server AAG.

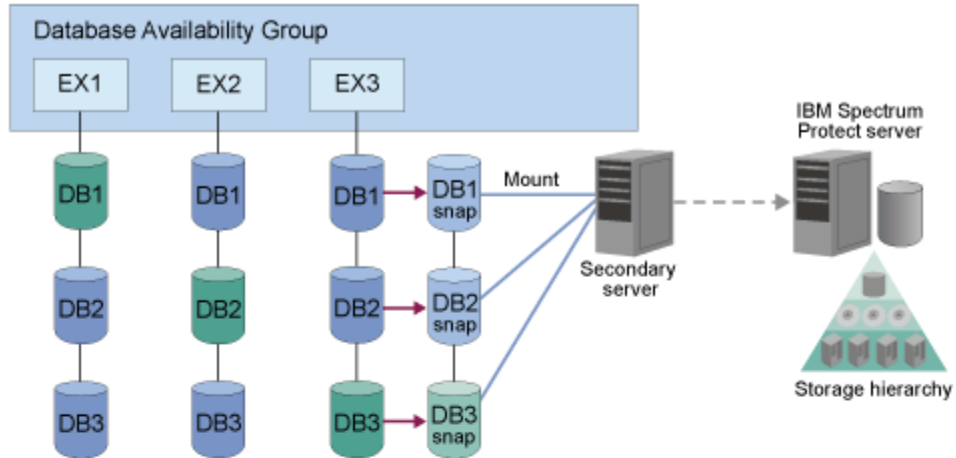


Figure 2. Example of how IBM Spectrum Protect Snapshot distributes VSS backups

In this deployment, you can install a Data Protection client component on one DAG node or AAG replica where the storage snapshot backup is to be completed. You can also define a CMD type schedule to run a CMD file with a backup command similar to the following example:

```
tdpexccc backup * full /backupmethod=VSS /backupdest=Both /offload
```

When you back up data to IBM Spectrum Protect from the VSS snapshot, use a passive copy. You can back up a passive copy, for example, the DB1 that is attached to EX3, so that the primary copy is not affected. You can also complete an offloaded backup by using the passive copy.

Similarly, as the illustration suggests, you can do both: back up data to IBM Spectrum Protect from the VSS snapshot by using a passive copy, and complete an offloaded backup by using the passive copy.

Examples of IBM SAN Volume Controller and IBM Storwize V7000 configuration scenarios

Configuration examples are scenarios that you can use to help you plan your data backup and recovery solutions.

Production application data is on standard volumes. Keep 14 snapshot backup versions. Use minimum storage space for snapshot backup versions. A full physical copy is not required. Complete two VSS backups per day.

SAN Volume Controller and Storwize V7000 settings

Create 14 space-efficient target volumes for each source volume to be protected. Enable the autoexpand option for the space-efficient target volumes. Add the space-efficient target volumes to the VSS_FREE pool.

VSS Provider settings

Set the background copy rate to 0.

IBM Spectrum Protect Snapshot settings

Set the policy to retain 14 local backup versions. Schedule snapshot backups as required by setting the backup destination option to LOCAL.

After 14 VSS backups are completed, the 15th VSS backup causes the oldest backup to be deleted and reuses that target set.

Production application data is on standard volumes. Keep one snapshot backup version. Use minimum storage space for snapshot backup versions. A full physical copy is not required. Perform one VSS backup per day and send the backup to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create two space-efficient target volumes for each source volume to be protected. Enable the autoexpand option for the space-efficient target volumes. Add the space-efficient target volumes to the VSS_FREE pool.

VSS Provider settings

Set the background copy rate to 0.

IBM Spectrum Protect Snapshot settings

Set the policy to retain two local backup versions. Schedule snapshot backups as required by setting the backup destination to BOTH

Set the policy for local snapshot backups to retain n+1 backup versions so that n snapshot backups are available for restore. Otherwise, a local backup version might not be available if a VSS backup fails after the prior backup was deleted.

Production application data is on standard volumes. Keep one snapshot backup version. A full physical copy is required. Minimize space usage of background copies. Perform one VSS backup per day and send the backup to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create one standard target volume for each source volume to be protected. Add standard target volumes to the VSS_FREE pool.

VSS Provider settings

Use the default background copy rate of 50. Configure a custom value to use incremental FlashCopy.

IBM Spectrum Protect Snapshot settings

Set the policy to retain one local backup version. Schedule snapshot backups as required by setting the backup destination to BOTH.

When you use incremental FlashCopy backup processing, the VSS provider does not delete the single snapshot target set even though IBM Spectrum Protect Snapshot software deletes the prior VSS snapshot before it creates a new snapshot.

Production application data is on standard volumes. Keep two snapshot backup versions. Full physical copies are required for local backup versions. Begin VSS backups every 12 hours with one backup sent to IBM Spectrum Protect daily.

SAN Volume Controller and Storwize V7000 settings

Create three standard target volumes for each source volume to be protected. Add standard target volumes to the VSS_FREE pool .

VSS Provider settings

Use the default background copy rate of 50.

IBM Spectrum Protect Snapshot settings

Set the policy to retain three local backup versions. Schedule VSS backups as follows: set the backup destination to LOCAL at 11 : 00, set the backup destination to BOTH at 23 : 00.

Set the policy for local snapshot backups to retain n+1 backup versions so that you can restore n snapshot backups.

Production application data is on standard volumes. Keep four snapshot backup versions. Use minimum storage space for snapshot backup versions. A full physical copy is not required. Perform VSS backups every six hours with one backup daily sent to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create five space-efficient target volumes for each source volume to be protected. Enable the autoexpand option for the space-efficient target volumes. Add space-efficient target volumes to the VSS_FREE pool .

VSS Provider settings

Use the default background copy rate of 0.

IBM Spectrum Protect Snapshot settings

Set the policy for local snapshot backups to retain five local backup versions. Schedule VSS backups as follows: set the backup destination to LOCAL at 06 : 00, 12 : 00, and 18 : 00, set the backup destination to BOTH at 00 : 00.

- Set policy to retain n+1 backup versions so that n snapshot backups are available for restore

Production application data is on space-efficient volumes. Keep two snapshot backup versions. A full physical copy is required for local backup versions. Perform VSS backups every six hours with one backup daily sent to IBM Spectrum Protect.

SAN Volume Controller and Storwize V7000 settings

Create three space-efficient target volumes for each source volume to be protected. Allocate the same percentage of real storage as for source volumes. Add space-efficient target volumes to the VSS_FREE pool .

VSS Provider settings

Use the default background copy rate of 50.

IBM Spectrum Protect Snapshot settings

Set the policy to retain three local backup versions. Schedule VSS backups as follows: set the backup destination to LOCAL at 06 : 00, 12 : 00, and 18 : 00, set the backup destination to BOTH at 00 : 00.

Set the policy for local snapshot backups to retain n+1 backup versions so that n snapshot backups are available for restore operations. This setting allows thin provisioning for both source and target volumes, and allows them to grow together.

Chapter 5. Protecting your data

By using Microsoft Management Console (MMC) or the command-line interface, you can back up and restore Microsoft Exchange Server data, Microsoft SQL Server data, or custom application and file system data.

About this task

If required, you can manage your installations remotely.

Starting Microsoft Management Console

After you complete the configuration process, start Microsoft Management Console (MMC) to protect your Exchange or SQL Server data.

Before you begin

If you try to use IBM Spectrum Protect Snapshot for Microsoft Exchange Server or IBM Spectrum Protect Snapshot for Microsoft SQL Server before you complete the configuration process, the software does not function correctly.

About this task

IBM Spectrum Protect Snapshot for Microsoft Exchange Server and IBM Spectrum Protect Snapshot for Microsoft SQL Server software is displayed in MMC as a plug-in. MMC uses a navigation tree to organize the computer data that is registered. Each computer icon that is followed by the word *Dashboard* represents a physical computer.

When you register a computer, information about the computer is collected and stored. Password information is encrypted and stored separately. The computers that are registered are tracked with a globally unique identifier (GUID). The GUID is used when you back up and restore data.

You can create groups of computers. These groups consolidate information when you view the dashboard, prepare reports, and run group commands. By default, the computers in a group are selected when you complete tasks for the group, but you can select more computers in the tree to include in an operation.

Procedure

- To start MMC, click **Start > All Programs > IBM Spectrum Protect Snapshot > IBM Spectrum Protect Snapshot Management Console**.

Starting the IBM Spectrum Protect Snapshot command-line interface

You can start the IBM Spectrum Protect Snapshot for Exchange Server or IBM Spectrum Protect Snapshot for SQL Server command-line interface by using a Windows command prompt with administrative privileges. Alternatively, you can start the command-line interface from Microsoft Management Console (MMC).

Procedure

1. Start MMC.
2. In the navigation tree, select the computer node where you want to run the commands.
3. Expand the **Protect and Recover Data** node.
4. In the navigation tree, select an Exchange Server node.

- Click the **Automate** tab. An integrated command line is available in the task window. You can use the interface to enter PowerShell cmdlets or command-line interface commands. The output is displayed in the main window.
- From the drop-down list, change **PowerShell** to **Command Line**.

Getting help for IBM Spectrum Protect Snapshot commands

By issuing the **help** command at the command prompt, you can view a complete list of commands and associated parameters for IBM Spectrum Protect Snapshot for Microsoft Exchange Server, IBM Spectrum Protect Snapshot for Microsoft SQL Server, and IBM Spectrum Protect Snapshot for File Systems and Custom Applications.

Procedure

- Use the following methods at the command prompt.

<i>Table 15. IBM Spectrum Protect Snapshot help commands</i>	
If you are using:	Issue this command
IBM Spectrum Protect Snapshot for Microsoft SQL Server	<p>tdpsqlc <i>?command_name</i> where <i>command_name</i> is the name of the IBM Spectrum Protect Snapshot command.</p> <p>For example:</p> <pre>tdpsqlc ? restore full</pre>
IBM Spectrum Protect Snapshot for Microsoft Exchange Server	<p>tdpexcc <i>?command_name</i> where <i>command_name</i> is the name of the IBM Spectrum Protect Snapshot command.</p> <p>For example:</p> <pre>tdpexcc ? backup</pre>
IBM Spectrum Protect Snapshot for File Systems and Custom Applications	<p>fcmcli <i>?command_name</i> where <i>command_name</i> is the name of the IBM Spectrum Protect Snapshot command.</p> <p>For example:</p> <pre>fcmcli ? backup</pre>

Determining managed storage capacity

You can track the capacity of managed storage assets. This information can be useful when you are calculating storage requirements for license renewal.

About this task

Typically, the capacity that is used by server data differs from the capacity of the volume that contains that data. For example, a set of databases might require a capacity of 1 GB and be on a 10 GB volume. When a snapshot of the volume is created, the IBM Spectrum Protect Snapshot managed capacity measurement is 10 GB.

Procedure

- From Microsoft Management Console (MMC), select an Exchange Server, SQL Server, or file system instance.

2. On the **Protect**, **Recover**, or **Automate** tab, in the **Actions** pane, click **Properties**.
3. Select **Managed Capacity** from the list of available property pages.
The managed capacity is calculated and displayed.
4. To view a list of the volumes that contain backups and their respective managed capacities, click **Show Details**.

Protecting Microsoft Exchange Server data

With IBM Spectrum Protect Snapshot for Exchange Server, you can back up and restore Microsoft Exchange Server databases.

About this task

You can create point-in-time snapshots of a Microsoft Exchange Server.

Note: For information about protecting Microsoft Exchange Server data in VMware environments, see chapter *Protection for in-guest applications* in the *IBM Spectrum Protect for Virtual Environments: Data Protection for VMware User's guide*.

Related tasks

[Configuring IBM Spectrum Protect Snapshot to restore mailboxes from mounted Exchange Server database files](#)

The **Mailbox Restore Only** configuration wizard is useful when you do not have to configure additional IBM Spectrum Protect Snapshot software functions.

Prerequisites

With IBM Spectrum Protect Snapshot, you can back up and restore Exchange Server data and protect your Exchange Server environment.

You can use IBM Spectrum Protect Snapshot for Microsoft Exchange Server to run backup and restore operations in a Database Availability Group (DAG) environment. A DAG consists of mailbox servers that provide recovery from database, server, or network failures. DAGs provide continuous replication and continuous mailbox availability.

Security requirements for backup and restore operations

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

If your user name is authorized by the security policy in your organization, you can add user names in the Exchange Organization Management role group or subgroups. Users whose name is in the Exchange Organization Management role group or subgroups can complete mailbox restore operations. Users whose name is not in the Exchange Organization Management role group or subgroups might experience slower performance when completing restore operations.

You must define a minimum set of management roles and role scope for the Exchange user.

- Set the role and scope:

Management roles

"Active Directory Permissions", "Databases", "Disaster Recovery", "Mailbox Import Export", "View-Only Configuration", and "View-Only Recipients".

To restore an Exchange 2013 or later public folder mailbox, the Exchange user must also have the **Public Folders** management role. To restore mail to a Unicode PST file, the Exchange user must have the **Mailbox Import Export** management role.

The following Exchange Powershell cmdlet sets RBAC permissions:

```
New-RoleGroup -Name "My Admins" -Roles "Active Directory Permissions",
"Databases",
```

"Disaster Recovery", "Mailbox Import Export", "Public Folders",

"View-Only Configuration", "View-Only Recipients" -Members operator1

The preceding example creates a new group, My Admins, with minimum roles to run IBM Spectrum Protect Snapshot, and assigns user operator1 to this group. The operator1 user can run IBM Spectrum Protect Snapshot but with limited Exchange privileges, for example, the user cannot create or remove a user mailbox.

Management role scope

Ensure that the following Exchange objects are in the management role scope for the user name who is logged on to the Exchange Server:

- The Exchange Server that contains the required data
 - The recovery database that IBM Spectrum Protect Snapshot creates
 - The database that contains the active mailbox
 - The database that contains the active mailbox of the user who completes the restore operation
- Verify that the Exchange user name is a member of a local Administrator group, and has an active Exchange mailbox in the domain.

By default, Windows adds the Exchange Organization Administrators group to other security groups, including the local Administrators group. For Exchange users who are not members of the Exchange Organization Management group, you must manually add the user account to the local Administrators group. By using the Local Users and Groups tool on the computer of the domain member, select **Administrative tools > Computer Management > Local Users and Groups tool**. On a domain controller computer that does not have a local Administrators group or Local Users and Groups tool, manually add the user account to the Administrators group in the domain by selecting **Administrative tools > Active Directory Users and Computers tool**.

Related information

Software requirements for backup and restore operations

To protect Microsoft Exchange Server 2013, 2016, and 2019 data, verify that your environment is set up correctly.

Ensure that your environment is set up to meet the following requirements.

Microsoft Exchange Server requirements

IBM Spectrum Protect Snapshot for Microsoft Exchange Server requires that you have local Administrator privileges.

Membership in the Organization Management group is not required because you might not want to grant Organization Management group permissions to all Exchange Server backup and restore operators. Instead, you can define customized role-based access control (RBAC) roles and management role scope so that Exchange Server users can run only limited operations within a limited scope.

Microsoft Exchange Server 2013 requirements

In Exchange Server 2013 mailbox restore operations, the MAPI clients must use the Remote Procedure Call over HTTP protocol (RPC over HTTPS, also known as Outlook Anywhere). You cannot use the RPC over TCP because Microsoft does not use that protocol.

Use Exchange Server 2013 CU2 or later versions, and download the correct MAPI. These software requirements are documented in the Hardware and Software Requirements technote at this location: [All Requirements\(http://www.ibm.com/support/docview.wss?uid=swg21219345\)](http://www.ibm.com/support/docview.wss?uid=swg21219345). Follow the link to the requirements technote for your specific release or update level.

After you configure your environment, mailbox restore operations work in the same way as with previous versions of Microsoft Exchange Server.

Related concepts

Security requirements for backup and restore operations

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Software requirements for mailbox restore operations

When you restore mailboxes and mailbox data, you can choose where to restore the mail and how to restore the mail. You can restore mailbox data from the GUI or command-line interface.

From these interfaces, you can restore data interactively by using the Mailbox Restore Browser or directly from Exchange Server database files. When you restore mailboxes and mailbox data on Exchange Server 2013 or later, ensure that your environment is set up to meet the following requirements:

- Ensure that the administrator account that is used to perform the mailbox restore operation has an active Exchange mailbox in the domain.
- Ensure that the user name who is logged in has role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- Ensure that the directory where you restore a mailbox has enough temporary disk space to store the entire mailbox database and log files. To specify the restore directory path, use the following settings on the **General property** page for the Exchange Server workload:

- **Temporary Log Restore Path**

- **Temporary Database Restore Path**

If you do not specify a directory, the database files are restored into a directory that is specified by the TEMP environment variable.

- **Exchange Server 2013:** Ensure that the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects is installed on the Exchange Server that you use to run the mailbox restore operations. The correct version is identified in the Hardware and Software Requirements technote that is associated with the level of your software. This technote is available at this web page: [All Requirements](#)

Follow the link to the requirements technote for your specific release or update level.

- **Exchange Server 2016 or later:** Install Microsoft 32-bit Outlook 2016 or later version as the MAPI client on the same server that Data Protection for Exchange Server uses for mailbox restore operations.

The amount of time that is needed to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

Related concepts

Exchange mailbox restore operations

By using IBM Spectrum Protect Snapshot, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

Microsoft Exchange Server backup methods

Depending on your Exchange Server environment, you can run only full backups, full plus incremental backups, or full plus differential backups. Your backup strategy might also include backing up data to IBM Spectrum Protect or local shadow volumes.

As you consider Exchange Server backup methods, understand all aspects of Exchange Server disaster recovery and the backup recommendations that Microsoft provides.

Follow these guidelines:

- Do not implement incremental and differential backups together.
- If you choose a strategy that involves incremental or differential backups, you must disable circular logging on the databases of the Exchange Server.

Full backup method

Use the full backup method during low usage times because a full backup can take a long time to run. However, the restore process is the most efficient because only the most recent full backup is restored.

Full backup plus incremental backup method

Use the full backup plus incremental backup method when the normal backup schedule or network capacity cannot support a full backup.

To minimize the effect on the backup schedule and network traffic during peak times, you can run a periodic full backup, followed by a series of incremental backups. For example, you can schedule full backups on the weekend and incremental backups during the week. You can run full backups during low usage times and when increased network traffic can be tolerated.

If you use this backup strategy, modify the IBM Spectrum Protect storage management policies to ensure that all incremental backups are collocated on the IBM Spectrum Protect server. In this way, you can improve data restore performance by reducing the number of media mounts that are necessary to restore a series of incremental backups.

Full backup plus differential backup method

Use the full backup plus differential backup method if your backup schedule and network capacity can facilitate backing up all transaction logs that accumulate between full backup operations. This strategy requires that only one differential backup plus the last full backup be transferred to complete a restore operation. However, the same amount of data must be transferred in the differential image, as in the series of incremental backup operations.

Therefore, a full backup plus differential backup policy increases network traffic and IBM Spectrum Protect storage usage. This policy assumes that the differential backups are processed as often as the incremental backups.

Consider the potential advantages and whether you can justify the additional resources that are necessary to resend all prior transaction logs with each subsequent differential backup.

IBM Spectrum Protect backups versus local shadow volumes backups

When you create a policy for your backups, you must choose whether to back up data to IBM Spectrum Protect storage versus VSS disks. Data backups to IBM Spectrum Protect typically takes longer to process than backups to local shadow volumes.

Backing up Exchange Server data to IBM Spectrum Protect is necessary when long-term storage is required. For example, saving Exchange Server data on tape for archival purposes requires long-term storage. IBM Spectrum Protect backups are also necessary for disaster recovery situations when the disks that are used for local backups are unavailable.

By maintaining multiple backup copies on IBM Spectrum Protect server storage, a point-in-time copy is available if backups on the local shadow volumes become corrupted or deleted.

Local shadow volumes

When you back up data to local shadow volumes, ensure that sufficient local storage space is assigned to the local shadow volumes. Create different sets of policies for backups to both local shadow volumes and to IBM Spectrum Protect server storage. If you use a VSS provider other than the Windows VSS System Provider, follow the backup recommendations of the VSS provider.

You can run backups to local shadow volumes by time and backup versions. It is more effective to base policy for local backups on version limits because local snapshots are created more frequently and VSS storage provisioning and space limitations apply. In Database Availability Group (DAG) environments, all of the DAG members must use the same local VSS policy.

Environment and storage resources also impact how many backup versions you can maintain on local shadow volumes for VSS fast restore and VSS instant restore operations, and on IBM Spectrum Protect server for VSS restore operations.

Database Availability Group backup and restore operations

To optimize use of available server resources, Database Availability Group (DAG) members often store a subset of the Exchange Server databases in a combination of active and passive copies.

Typical DAG configuration

In the following example, three copies of five databases span five servers in a DAG. This configuration ensures that two servers in the DAG never have the same set of database copies. The configuration also provides greater resilience to failures. Specifically, three servers must fail before the servers lose access to a database.




Figure 3. Typical DAG configuration

Typical data protection deployments in DAG environments

You can back up data from any DAG member and restore the data to any DAG member. You can also back up data from either the active or passive copy. Full and incremental database backups do not have to be completed from the same DAG member. All databases included in a VSS type backup are integrated.

The following figure illustrates a deployment of a backup task that is distributed across DAG members.

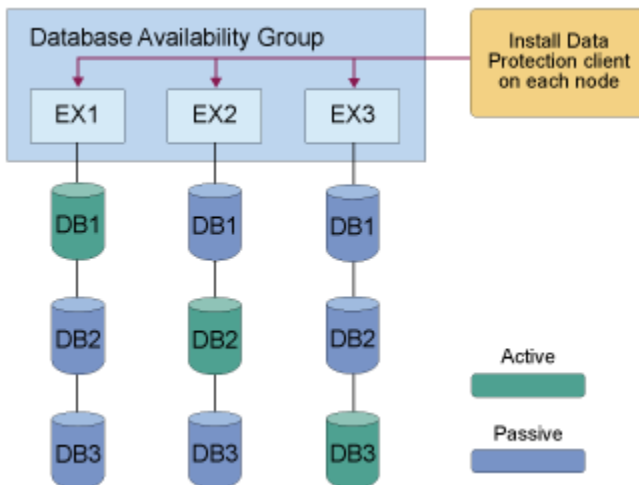


Figure 4. Example of backups that are distributed across DAG members

To specify a backup of all DAG nodes, issue the same backup command on each node. The command file contains separate backup commands per database. For example:

```

tdpexcc backup DB1 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB2 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB3 full /minimumbackupinterval=60 /preferdagpassive

```

In this deployment, one schedule applies to all nodes. The same backup command file is used for each node.

The following figure illustrates another possible backup task distribution across DAG members.

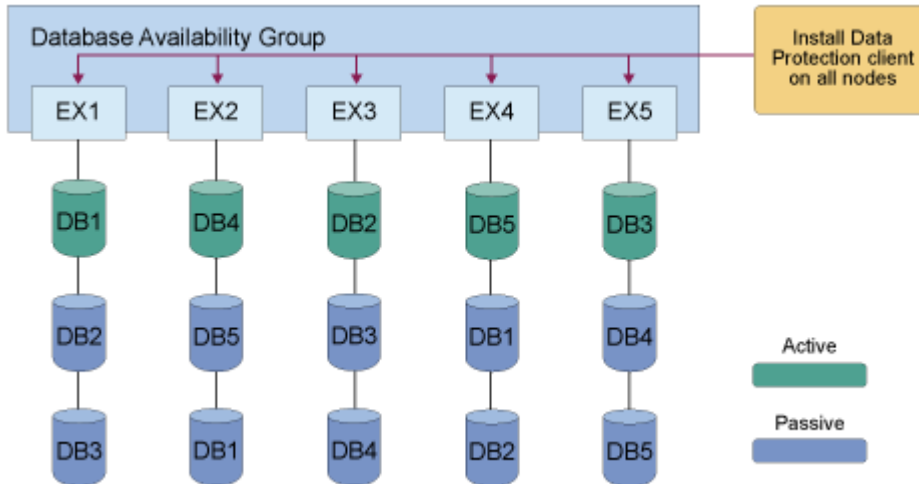


Figure 5. Another example of backups distributed across DAG members

In this deployment, one schedule applies to all nodes. The same backup command file is used for for all nodes. The command file contains separate backup commands per database on that node. For example:

```

tdpexcc backup DB1 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB2 full /minimumbackupinterval=60 /preferdagpassive
tdpexcc backup DB3 full /minimumbackupinterval=60 /preferdagpassive

```

Best practices for backing up a Database Availability Group

When you back up data, distribute the backup workload for scalability and isolate backup activity to a dedicated backup node. When you isolate backup activity, it minimizes the impact to production databases.

As a best practice, identify all replica copies of the same database and eliminate redundant backups of the same databases. You can apply retention policies to databases. Back up databases from any node in the availability group and run restore operations from any node in the availability group.

Complete backups for replicated database copies from the same Exchange Server. Additionally, complete backups on the passive database copies. When you backup passive database copies, you do not increase the load on the production Exchange Server.

When you back up databases, follow these guidelines:

- Use a DAG member to store DAG database backups.
- Ensure that the same VSS policy applies to all DAG members.
- Ensure that the first backup is a FULL backup when you move backups to DAG member backups.
- Ensure that previous backups are manually deleted after you move backups to DAG member backups, assuming that those backups are no longer needed.
- Run backups from a passive database copy to avoid increasing the load on the active databases.
- Schedule all DAG members that have a copy of the database to back up the database at the same time. To set the minimum amount of time before a backup of another DAG copy of the same database is allowed, specify the **MINIMUMBACKUPINTERVAL** parameter. When you specify this parameter, only one backup is taken per backup cycle.

- If the Exchange Server database belongs to a DAG and is an active database copy, specify the **/EXCLUDEAGACTIVE** parameter to exclude the databases from the backup.
- If the Exchange Server database belongs to a DAG and is a passive database copy, specify the **/EXCLUDEDAGPASSIVE** parameter to exclude the databases from the backup.
- If the Exchange Server database does not belong to a DAG, specify the **/EXCLUDENONDAGDBS** parameter to exclude the databases from the backup.
- To a backup is to be taken from a passive copy unless no valid passive copy is available, specify the **/PREFERDAGPASSIVE** parameter.
- To bypass an integrity check if two or more valid database copies exist in a DAG, specify the **/SKIPINTEGRITYCHECK** parameter.

Best practices for restoring a Database Availability Group

In a DAG environment, you must restore databases on an active database copy. To restore to a passive database copy, you must first move the copy to the active state. After the restore operation is complete, you can move the active database copy to the passive state.

If you back up data to a local system, you can complete data restore operations only on the Exchange Server where the backup is taken.

Ensuring successful MAPI connections

If you use Exchange Server 2013, use the **MAPI Settings** property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables IBM Spectrum Protect Snapshot to connect to the Exchange Server.

Before you begin

Ensure that the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects is installed on the Exchange Server. The correct version is identified in the Hardware and Software Requirements technote that is associated with the level of your software.

About this task

For mailbox restore operations to succeed in Exchange Server 2013 environments, the MAPI client must use Remote Procedure Call over HTTPS (RPC over HTTPS), also known as Outlook Anywhere. You cannot use RPC over TCP.

Procedure

1. From Microsoft Management Console (MMC), select an Exchange Server instance.
2. On the **Protect** tab, click **Properties** in the Action pane.
3. Select **MAPI Settings** from the list of property pages.
4. Verify that the following information is correct in the Exchange Server environment:
 - The **mailbox alias** field shows the mailbox that you are logged in to. Verify that you can open the mailbox in Microsoft Outlook or Outlook Web Access (OWA).
 - The **Exchange Profile Server** field shows the correct mailbox endpoint on the Exchange Server that has the Client Access Server (CAS) role. Verify that you can open the target mailbox in Outlook or OWA.
5. Edit the registry key only if the default value is incorrect. Use one of the following methods.
 - Enter the registry key value in the `RpcHttpProxyMap_TSM` field.
 - Enter the `Domain` field value and select or clear the **Use HTTPS authentication** check box. When you change either of these values, the values of the registry key automatically updates in the `RpcHttpProxyMap_TSM` field.

Consider that the values that you enter override the registry key that is in the `HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem` directory. If you modify the registry incorrectly, the connection to the Exchange Server might fail.

RpcHttpProxyMap_TSM

Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment. By default, the following format is used.

```
Domain=Proxy Server,RpcHttpAuthenticationMethod,  
RpcAuthenticationMethod,IgnoreSslCert
```

For example:

```
companyname.local=https://exchange.companyname.com,ntlm,ntlm,false
```

where:

- *Domain* value is the domain suffix of the personalized server ID, for example, `companyname.local`. Specify any domain or a substring of a domain, or the asterisk (*) and question mark (?) wildcard characters, for example, `*.companyname.local`.
- *Proxy Server* value is the RPC proxy server that has the Client Access Server (CAS) role. Specify the fully qualified domain name (FQDN) of the RPC proxy server. Precede the FQDN by `http://` for an HTTP connection, or `https://` for an HTTPS connection. For example, `https://exchange.companyname.com`
- *RpcHttpAuthenticationMethod* value is the method that is used to authenticate RPC-over-HTTP connections. Specify NTLM, Basic, Negotiate, or WinNT.
- *RpcAuthenticationMethod* value is the method that is used to authenticate RPC-over-TCP connections. Specify NTLM, Negotiate, WinNT, Anonymous, or None.
- *IgnoreSslCert* value indicates whether the Exchange Server validates SSL certificates. For the Exchange Server to ignore invalid certificates, specify `False`.

Domain

Change the domain name to reflect the correct domain if for example, you have multiple domains, or the default domain value is incorrect. To match all domains, enter the asterisk (*) wildcard character. When you change this domain value, the *Domain* value of the registry key automatically updates in the `RpcHttpProxyMap_TSM` field.

Use HTTPS authentication

Select this check box if RPC-over-HTTPS is enabled for the Exchange Server that is hosting the MAPI profile. Otherwise, clear this check box to ensure that HTTP authentication is used for RPC-over-HTTP connections. When you change this authentication value, the *RpcAuthenticationMethod* value of the registry key automatically updates in the `RpcHttpProxyMap_TSM` field.

Related tasks

[Troubleshooting mailbox permissions, authentication methods, and registry key settings in a Microsoft Exchange 2013 environment](#)

To resolve mailbox restore errors in an Exchange Server 2013 environment, ensure that the Exchange Server mailbox permissions, authentication methods, registry key settings, and the Client Access Server (CAS) role are configured correctly.

[Troubleshooting MAPI connection issues](#)

Related reference

[MAPI Settings](#)

If you use Exchange Server 2013, use the **MAPI Settings** property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables IBM Spectrum Protect Snapshot to connect to the Exchange Server.

Backing up Exchange Server data by using VSS

By using Microsoft Volume Shadow Copy Service (VSS), you can back up Exchange Server data and mount the backup if required.

Before you begin

- You must have a VSS provider that is configured for your environment.
- If you back up Exchange Server databases in a DAG environment, and you want to back up your databases to a common node, ensure that you set up a DAG node name (DAGNODE).

Tip: Backing up DAG databases to a common node is helpful when you want to manage backups with a single policy, regardless of which DAG server completes the backup.

You can set up the DAG node name in the **DAG Node** field in the **TSM Node Names** page of the IBM Spectrum Protect configuration wizard, or in the **Back up DAG databases to common node** field in the **General** properties page for your Exchange Server workload.

- Do not mix persistent and nonpersistent VSS snapshots.

Procedure

1. Start Microsoft Management Console (MMC) and click **Exchange Server** in the navigation tree.
2. On the **Protect** tab, select one or more databases to back up. Alternatively, click the **Protect Data** shortcut in the start page of MMC.
 - a) Filter the list of available databases in the results pane by entering a keyword in the **Search** field.
 - b) If you are running backup operations in an Exchange Server DAG environment, you can back up an active database copy or passive database copy. View the copy status in the **DAG Status** column on the **Protect** tab.
3. Specify the backup options. If the backup options are not displayed, click **Show Backup Options**.
 - To use offloaded backups, set the **Offload** option to **True**.
If you use offloaded backups, specify the remote client node, **RemoteDSMAGENTNode**, that runs the VSS offloaded backups on a remote computer. This option applies only to the IBM Spectrum Protect configuration.
 - Select **Skip Integrity Check** and choose one of the following options.

Task	Action
Bypass integrity checking for all database and log files	Select Yes
Run integrity checking to verify that all database and log files are free of errors	Select No This option is the default.
Bypass integrity checking for database files only if at least two valid copies of a database (one active and one passive copy) exist in a DAG	Select Skip Database Check If Healthy
Bypass integrity checking for database and log files only if at least two valid copies of a database (one active and one passive copy) exist in a DAG	Select Skip Database And Log Check If Healthy

- If you are scheduling the backup of databases in an Exchange Server DAG, set the **Minimum Backup Interval** value to the minimum amount of time, in minutes, before a backup of another copy of the same DAG database can begin. The default value is 0, which means that you can back up the database again immediately after a backup operation of that database is complete. The time of the last database backup is determined from the Exchange Server and not the IBM Spectrum Protect server.

This option specifies that only one database copy can be backed up within a time frame. This option prevents all members in a DAG from backing up the database. Specify this setting for tasks that are scheduled to run when you click **Run Scheduled**.

- If you are scheduling the backup of databases in an Exchange Server DAG, set **PreferDAGPassive** option to **True** to skip the backup for an active database copy unless no valid passive copy is available. If no valid passive copy is available, the backup is created from the valid active database copy.

Specify this setting for tasks that are scheduled to run when you click **Run Scheduled**.

4. Optional: Choose a mode for the current task:

- **Run Interactively:** Click this item to run the current task interactively. This selection is the default.
- **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.

5. To create the backup, select the backup action in the **Actions** pane.

You can run a full, copy, incremental, or differential backup with the VSS backup method.

Related concepts

Offloaded VSS backups

By running an offloaded backup, you can move the backup load from the production system to another system. You can reduce the load on network, I/O, and processor resources during backup processing.

Related tasks

[“Restoring a Database Availability Group database backup ” on page 106](#)

You can restore a replicated database copy in a Database Availability Group (DAG).

Mounting Exchange Server backups

To see a copy of Exchange Server data from a specific point in time, mount a snapshot backup.

About this task

A copy of data from a specific time is also known as a point-in-time consistent copy or online snapshot.

Restriction: You cannot use Microsoft Management Console (MMC) to mount a backup to a different server. To mount a VSS snapshot to a remote server, enter the **mount backup** command at the command line, or use the **Mount -DpExcBackup** cmdlet.

When you submit a mount request, all of the volumes that are contained in the original snapshot set are imported. If the number of volumes that are imported exceed the maximum number of allowable mapped volumes for the environment, the mount operation can fail.

You can mount VSS snapshot backups either as read-only or read/write. When a snapshot backup is mounted as read/write, you can do individual mailbox or mail item restores without needing to copy the Exchange database file from the snapshot backup into the recovery database (RDB); which greatly reduces the restore time. There are two variations of the mount read/write option:

- **Mount read/write (modifies backup, applies to COPY backups only)**

For VSS providers that support transportable shadow copies, you can mount a COPY type backup as read/write. After mounting, your COPY backup is marked as modified and while you can mount it again in the future, this backup can no longer be used as a restore point in future database restore operations. It can be used for mailbox restore operations only. All databases on the snapshot volume that are mounted as read/write are marked as modified.

- **Mount read/write (without modifying backup)**

This mount option is only available for the following devices.

- SAN Volume Controller (SVC) devices, which requires IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later.
- XIV system devices, which requires IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

With these options, you can mount writable VSS snapshots of FULL or COPY backups.

Note: You can override your default mount options as specified in the configuration file by using either the **/MOUNTRW** parameter on the **mount backup** command, or the **Mount-DpExcBackup** cmdlet.

Procedure

1. Start MMC.
2. Click **Recover Data** in the welcome page of MMC.
3. In the **Actions** pane on the **Recover** tab, click **Mount Backup**.
4. Either type the path to the empty NTFS or ReFS folder where you want to mount the backup or browse to find the path. Click **OK**.
On the **Recover** tab, the backup that you mounted is displayed.
5. Optional: If required, select the **Mount the snapshots in read/write mode** option.
6. In the **Actions** pane, select the **Explore** and **Unmount Backup** tasks for the backup that you mounted.

Related reference

[Mount backup command](#)

To mount backups that are managed by IBM Spectrum Protect Snapshot for Exchange Server, use the **mount backup** command.

Deleting Exchange Server backups

You can remove an Exchange Server backup that you created with the VSS backup method. Use this procedure only for deletions that are outside the scope of your standard policy management deletions.

Before you begin

Typically, backups are deleted automatically based on user-defined policy management settings. This procedure is necessary only if you must delete backups that are outside the scope of IBM Spectrum Protect Snapshot policy definitions.

If you back up Exchange Server Database Availability Group (DAG) databases to IBM Spectrum Protect, you can delete the database backup from the DAG member to a local shadow volume only from the Exchange Server on which the backup is created.

If you delete a remotely-mounted backup, the snapshots and the relationship between the source and target volumes on the storage device are also deleted. However, the target volume that is imported and mounted might continue to exist. In addition, the target volume might not be available to the server where the remote mount occurred. The operations to the target volume depend on the VSS hardware provider and the storage device implementation.

After the maximum number of remotely-mounted backup versions or the maximum number of days to retain a backup is exceeded, the associated backup is expired and deleted.

Procedure

1. Start Microsoft Management Console (MMC).
2. Click **Recover Data** in the welcome page of MMC.
3. On the **Recover** tab for the Exchange Server instance, select **View: Database Restore**. In the **Results** pane, browse to and select one or more database backups to delete.
4. In the **Actions** pane, click **Delete Backup**.

While a backup is being deleted, two tasks are displayed in the task window to show that the deletion is in progress, and that the view is being refreshed.

Related tasks

Mounting VSS snapshots to remote servers

You can mount VSS snapshots to remote servers that other users can access. When you remotely mount a snapshot, you must use hardware that includes a feature for creating transportable snapshots. If you use a mix of hardware, all hardware must include a feature for creating transportable snapshots.

Setting data restore options in Microsoft Management Console

To optimize the data restore process for your environment, modify the default options that are available in Microsoft Management Console (MMC).

Procedure

1. On the **Recover** tab, select **Database Restore**.
2. Click **Show Restore Options** to modify the default restore options as follows:

<i>Table 17. Database restore options</i>	
Option	Action
Auto Select	<p>For this option, specify a value of Yes (default) to quickly select the backup objects to restore. With automatic selection, when you select the most recent backup to restore, all associated backups are automatically selected, up to the previous full backup. When you specify Yes, the automatic selection option applies to full backups, differential backups, and incremental backups, but not to copy backups. This option affects backups in the following ways:</p> <ul style="list-style-type: none"> • When you click a differential backup, the associated full backup is also selected. • When you click an incremental backup, the associated full backup and all associated earlier incremental backups are also selected. • For VSS backup, automatically selects all databases that were backed up together to the local destination. However, databases that were backed up to IBM Spectrum Protect are not automatically selected.
From Server	Enter the name of the server where the original backup is completed. The default value the local server.

Table 17. Database restore options (continued)

Option	Action
<p>Instant Restore</p>	<p>For this option, specify a value of Yes to use volume-level snapshot restore (instant restore) for local VSS backups if the backup exists on SAN-attached volumes. Specify a value of No to disable instant restore, which bypasses volume-level copy and uses file-level copy (fast restore) to restore the files from a local VSS backup. The default value is Yes, which uses volume-level snapshot restore if it is available.</p> <p>This option is available for VSS operations only. If you use instant restore for SAN Volume Controller earlier than version 5.1 or DS8000, ensure that any previous background copies that involve the volumes that are being restored are completed before you initiate the instant restore.</p> <p>This option is automatically set to No during <i>restore into</i> operations.</p> <p>In an instant restore operation, files on the destination file system are overwritten. Incremental and differential backups are automatically converted to file-level restores. An instant restore operation requires that the drive or volume where the mailbox database is located must be available. Any other process or application must not have access to the drive or volume.</p>
<p>Mount Databases After Restore</p>	<p>For this option, specify a value of Yes to automatically mount databases after backups are recovered. No is the default value for this option.</p>
<p>Run Database Recovery</p>	<p>If more backup files need to be restored before you run database recovery, select False.</p> <p>If you want to run database recovery after the restore is done, select True and specify whether only restored logs or both restored and current logs should be used for the recovery.</p> <ul style="list-style-type: none"> • Replay Restored and Current Logs Replays all transaction log entries both in the restored log files and in the log files on the server that have not been backed up. You cannot specify this option for instant restore operations. • Replay Restored Only Replays only restored logs. <p>Note: From the MMC, the ERASEexistinglogs parameter can be applied to erase the existing transaction log files for the database that is being restored before you restore it.</p>

Restoring an Exchange Server database

You can use the *restore into* function to restore an Exchange Server database backup to a recovery database or alternate database. You can also restore a DAG active or passive database copy to a recovery database or alternate database.

Before you begin

- Ensure that your system is set up to use the DAG node name (DAGNODE). You can specify the DAG node name in the **DAG Node** field in the IBM Spectrum Protect **Node Names** page of the IBM Spectrum Protect configuration wizard, or in the **Back up DAG databases to common node** field in the **General** properties page for your Exchange Server workload.
- You can restore mailboxes with the Mailbox Restore Browser or Mailbox Restore functions. In some rare cases, however, you might want to restore data into a recovery database or alternate database. Ensure that a recovery database or alternate database exists before you attempt the restore operation.

About this task

- For database backups in the Exchange Server Database Availability Group (DAG) environment, you can restore a database regardless of which DAG member the database was backed up from because all database copies are backed up by using a single DAG node. Local backups must be restored on the node where the backup was completed.

In a stand-alone environment, you cannot back up a database from one DAG member and restore it to a different DAG member. Backups of the same database are managed with the same policy, regardless of whether the database is active or passive at the time of the backup.

- Running any type of *restore into* function automatically disables VSS instant restore capability.

When you restore a database by using instant restore processing, data that exists in the destination database is overwritten, and is no longer available after restore processing is complete. When you restore a database by using the *restore into* function, you restore data to an alternate target destination. The data is not restored to the original source destination. For the restore operation to be successful, the alternate target destination must be of equal or greater size as the original source volume.

- To complete restore operations, backups must be taken on the same version of Exchange Server.
- You cannot use multiple instances of IBM Spectrum Protect Snapshot for Exchange Server to restore databases into the recovery database simultaneously.

Procedure

1. From Microsoft Management Console (MMC), click **Recover Data** in the welcome page.
2. On the **Recover** tab for the Exchange Server instance, select **View: Database Restore**.

In the **Results** pane, browse to the databases that are available to restore. The following options are available:

Table 18. Database restore selection options

Option	Action
<p>Filter</p>	<p>Use the filter options to narrow the list of databases in the result pane.</p> <ol style="list-style-type: none"> Click Show Filter Options and Add Row. In the Column Name field, click the down arrow and select an item to filter. <p>For database backups in the Exchange Server DAG environment, the Server column displays the name of the DAG and the server that created the backup in this format:</p> <pre>DAGNAME\SERVERNAME</pre> <p>where DAGNAME is the name of the DAG, and SERVERNAME is the name of the server (DAG member) that created the backup. For example:</p> <pre>TSMDAG4\AVOCADO</pre> <p>To filter by Backup Date, click the default date and time to edit the table cell. To change the date, click the arrow button that is displayed at the end of the cell. The calendar date selection tool is displayed. After you select a date, to display the date in the field, press Enter. To edit the time, enter the time by using the 12-hour clock time convention such as 2 p.m.</p> <p>When you click Select All, all rows that reflect the filter specifications are selected.</p> <ol style="list-style-type: none"> In the Operator field, select an operator. In the Value field, specify a filter value. If you want to filter on more items, click Add Row. Click Apply Filter.
<p>Backups</p>	<p>Select the database to restore. You can click Active Backups to show only active backups, or click All Backups to show both active and inactive backups.</p>
<p>Search</p>	<p>In the Search field, enter a keyword to filter the list of available databases.</p>
<p>Refresh</p>	<p>Click Refresh to update the view with your changes.</p>

If you applied a filter, the objects on the server that match the filter or search criteria are listed in the **Recover** tab. The status area indicates the number of items that match the criteria n of x displayed, where n equals the number of objects that match the filter criteria, and x is the number of objects that are retrieved from the server. For example, 5 of 20 displayed. If you specify refresh options to further narrow your results, and click **Refresh** again, the objects on the server that match

the filtered and refresh options are displayed. Each time that you click **Refresh**, another query is run against the IBM Spectrum Protect server.

3. On the **Recover** tab for the Exchange Server instance, select one or more backups to restore.
If the **Auto Select** option is set to **Yes** in the **Restore Options** view, more backups that are necessary to restore the most recent backup are selected for you. If you do not want the additional selections that are made for you, set **Auto Select** to **No**.
4. Verify the restore options. If the restore options are not displayed, click **Show Restore Options**.
5. Optional: Choose a mode for the current task:
 - **Run Interactively**: Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled**: Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
6. Start the restore operation:
 - To restore the backup, right-click that backup name and select **Restore**. Alternatively, in the **Actions** pane, click **Restore**.
 - To restore the backup into another location, right-click and select **Restore Into** to specify a target location for the restore operation. A dialog window opens where you can specify the destination database.

Select the name of a database into which a VSS backup is restored.

VSS instant restore is available only for full or copy type backups that are on the disk devices that support this type of restore operation. During the VSS instant restore operation, the drive or volume where the database is located must not be accessed by any other process or application.

Restoring a Database Availability Group database backup

You can restore a replicated database copy in a Database Availability Group (DAG).

About this task

You can use Exchange Management Shell commands, which are provided in parentheses.

Procedure

1. Make the database that you want to restore active (**Move-ActiveMailboxDatabase**).
2. Suspend replication of all passive copies of the database (**Suspend-MailboxDatabaseCopy**).
3. Unmount the active mailbox database (**Dismount-Database**).
4. If you are using VSS instant restore, and the **During Instant Restore, automatically stop and restart necessary Microsoft Exchange services** option is not selected in Microsoft Management Console (MMC), or the **STOPSERVICESONIR** parameter is set to NO at the command line, stop the following replication services on all copies of the database.
 - (DAG environments only) Exchange Replication Service
 - (Exchange Server 2013 or later only) Exchange Search Host Controller Service
5. Restore the database and logs by using the command line or MMC.

Restriction: The database must not be mounted automatically after the restore. If you use MMC, ensure that the **Mount Databases After Restore** option is set to **No** in the **Restore** pane. If you use the command line, set the **/mountdatabases** parameter to NO.

However, if the **During Instant Restore, automatically stop and restart necessary Microsoft Exchange services** option is selected, or the **STOPSERVICESONIR** parameter is set to YES, you can set the **Mount Databases After Restore** option to YES.

6. If the service is stopped, start the replication service before you mount the active mailbox database. Otherwise, the database mount fails (**Mount-Database**).

7. Verify the health of the database before you update or reseed to replicated database copies. (**Get-MailboxDatabaseCopyStatus**)
8. Update or reseed all replicas (**Update-MailboxDatabaseCopy**). By completing this step, you can help to avoid potential transaction log synchronization problems that might arise if replication is resumed directly.
9. Move the active database to the server that you want (**Move-ActiveMailboxDatabase**).

Restoring mailbox data

IBM Spectrum Protect Snapshot backs up mailbox data at the database level, and also restores individual mailbox items from the database backup.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations.

If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) server, ensure that you configure the SMTP server before you start a restore operation. To set the configuration in the Management Console, right-click **Dashboard** in the tree view and select **Properties**. From the **E-mail** property page, enter the SMTP server and port.

About this task

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a .pst file. When you restore a mailbox to the mailbox restore destination, IBM Spectrum Protect Snapshot automatically restores the mail items in the Recoverable Items folder.
 - You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the mail items in the folders.
 - The mail items that you can restore depends on whether the mailbox is enabled for mailbox restore operations.
 - You can restore the Recoverable Items content for a public folder mailbox but not for each public folder in the public folder mailbox.
 - You can exclude the mail items in the Recoverable Items folder in mailbox restore operations.
 - You cannot create a subfolder in the Recoverable Items folder in a mailbox.
- In Exchange Server 2013 or later versions, you can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
 - To restore an Exchange 2013 or later public folder mailbox, the Exchange user must have the Public Folders management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange Server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange Server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange Server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.

If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.

- You might restore to a different public folder mailbox than the original mailbox if, for example, the public folder is relocated after the time of the backup. Before you complete the public folder restore

operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.

- In Exchange Server 2013 or later versions, you can restore an archive mailbox or a part of the mailbox, for example, a specific folder. You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

If you enable a user mailbox to be archived, ensure that the user is logged on to that mailbox at least once before you complete a backup and restore operation on the mailbox.

- If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all the mailboxes are in the same recovery database.
- By default, IBM Spectrum Protect Snapshot restores the latest backup that is available for the specified mailbox.

The amount of time that it takes to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

Procedure

1. Start Microsoft Management Console (MMC) and select **Exchange Server** in the navigation tree.
2. On the **Recover** tab for the Exchange Server instance, select the **Mailbox Restore** view.
3. Select one or more mailboxes to restore.

Restriction: A list of all available user mailboxes in the domain is displayed, including those mailboxes that were not backed up. Mailboxes that are not backed up cannot be selected for a restore operation. Only mailboxes that are backed up can be restored.

If you restore mail to a Unicode personal folder (.pst) file, or you restore a mailbox that is deleted or re-created after the time of the backup, IBM Spectrum Protect Snapshot for Microsoft Exchange Server requires a temporary mailbox to store the mailbox messages. Create a temporary mailbox by setting the Alias of temporary mailbox option on the **Properties** page, under the **General** tab.

Tip: Ensure that the temporary mailbox that you create is on a database with enough disk storage capacity to accommodate all of the mailbox items that you are restoring.

4. Optional: By default, the entire mailbox is restored. You can use the **Item-Level Mailbox Filters** to identify individual messages to restore:
 - a) Click **Show Filter Options** and **Add Row**.
 - b) In the **Column Name** field, click the down arrow and select an item to filter.
 - You can filter public mailbox folders only by the **Folder Name** column.
 - You can filter Unicode .pst files only by **Backup Date**, **Folder Name**, and **All Content** filters.
 - You can filter by backup date, and click the default date and time to edit the table cell. To change the date, click the arrow at the end of the cell. The calendar date selection tool is displayed. After you select a date, to display the date in the field, press **Enter**. To edit the time, enter the time by using the 12-hour clock time convention such as 2 p.m.

When you specify a backup date, IBM Spectrum Protect Snapshot for Microsoft Exchange Server searches for a backup that corresponds to that exact date. If a backup with that exact date is not found, IBM Spectrum Protect Snapshot for Microsoft Exchange Server selects the first backup after that date.
 - c) In the **Operator** field, select an operator.
 - d) In the **Value** field, specify a filter value.
 - e) If you want to filter on more items, click **Add Row**.
5. Specify the restore options by clicking **Show Restore Options**.

<i>Table 19. Database restore options</i>	
Task	Action
Keep Recovery Database After Restore	Use this option to retain a recovery database after a mailbox restore operation is complete. The default value is No . If you set the value to Yes , IBM Spectrum Protect Snapshot for Microsoft Exchange Server automatically retains the recovery database after mailbox restore processing.
Mailbox	If the alias of the mailbox to restore is not displayed in the list of mailboxes, specify the alias. This option overrides any selected mailboxes.
Mailbox Original Location	Use this option only if the mailbox was deleted or re-created since the time of the selected backup, and mailbox history is disabled. Specify the Exchange Server and the database where the mailbox was at the time of the backup. Use the following format: server-name, db-name, for example, serv1, db1.
Mark Restored Messages As Unread	Use this option to automatically mark the mailbox messages as unread after the restore operation is completed. The default value is Yes .
Use Existing Recovery Database	Use this action to restore the mailbox from an existing recovery database. The default value is Yes . If you set the value to No and a recovery database is mounted on the server before you restore the mailbox, IBM Spectrum Protect Snapshot for Microsoft Exchange Server automatically removes the recovery database during mailbox restore processing.
Enter Mount Point or Directory for MountRW Mailbox Restore	Select this option to specify either a directory path or a mount point for a read/write mount of a local VSS snapshot backup, that you want to use for a mailbox restore. If you do not use read/write mounts, no entry is necessary. Ensure that Use Existing Recovery Database value is set to No .

6. To complete the restore operation, click one of the following **Restore** options.

<i>Table 20. Restore options</i>	
Task	Action
Restore Mail to Original Location	Select this action to restore mail items to their location at the time of the backup operation.

<i>Table 20. Restore options (continued)</i>	
Task	Action
Restore Mail to Alternate Location	<p>Select this action to restore the mail items to a different mailbox.</p> <p>Note: If deleted mail items or tasks are flagged in the Recoverable Items folder of a mailbox, the items are restored with the flag attribute to the Flagged Items and Tasks view in the target mailbox.</p>
Restore Mail to non-Unicode PST file (Exchange Server 2013 only)	<p>Select this action to restore the mail items to a non-Unicode personal folders (.pst) file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location. Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory.</p> <p>If the .pst file exists, the file is used. Otherwise, the file is created.</p> <p>Restriction: The contents of each folder cannot exceed 16,383 mail items.</p>
Restore Mail to Unicode PST file	<p>Select this action to restore the mail items to a Unicode .pst file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location.</p> <p>You can enter a standard path name (for example, c:\PST\mailbox.pst) or a UNC path (for example, \\server\c\$\PST\mailbox.pst). When you enter a standard path, the path is converted to a UNC path. If the UNC is a non-default UNC path, enter the UNC path directly.</p> <p>Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory. If the .pst file exists, the file is used. Otherwise, the file is created.</p>

Table 20. Restore options (continued)	
Task	Action
Restore Public Folder Mailbox	<p>Select this action to restore a public folder mailbox to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>
Restore Mail to Archive Mailbox	<p>This action applies to a primary mailbox or an archive mailbox. Select this action to restore all or part of either type of mailbox to the original archive mailbox or to an alternate archive mailbox.</p> <p>You can filter the archive mailbox and restore a specific mailbox folder. In the Folder to be restored field, enter the name of the folder in the archive mailbox that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>In the Target archive mailbox field, specify the archive mailbox destination that you want to restore to.</p>
Exclude recoverable mail items while restoring the mailbox	<p>Apply this action if you are restoring an online, public folder, or archive mailbox to an original mailbox, alternate mailbox, or to a Unicode .pst file.</p> <p>Specify a value of Yes to exclude the mail items in the Recoverable Items folder in mailbox restore operations. No is the default value.</p>

Related concepts

[Exchange mailbox restore operations](#)

By using IBM Spectrum Protect Snapshot, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

[Security requirements for backup and restore operations](#)

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Related tasks

[Troubleshooting a MAPI error that prevents multiple mailboxes restoring in a Microsoft Exchange 2013 environment](#)

When you restore multiple mailboxes on a server that is running Exchange Server 2013, the mailbox restore operation might partially fail and report a MAPI error.

[Setting data restore options in Microsoft Management Console](#)

To optimize the data restore process for your environment, modify the default options that are available in Microsoft Management Console (MMC).

[Deleting mailbox history information](#)

Mailbox history includes only the mailboxes from databases that are backed up. If you back up mailbox history with a version of IBM Spectrum Protect Snapshot earlier than version 4.1, you can manually delete the old mailbox history.

[Troubleshooting mailbox restore errors](#)

If you encounter a mailbox restore error, determine whether the problem is reproducible on other Exchange Servers.

Related reference

[Restoremailbox command](#)

To restore mailbox-level data or mailbox-item-level data from IBM Spectrum Protect Snapshot for Exchange Server backups, use the **restoremailbox** command.

Restoring mailbox messages interactively with the Mailbox Restore Browser

You can use the **Mailbox Restore Browser** to interactively restore a mailbox or items from a mailbox on an Exchange Server.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations.

If you plan to restore mail or folders by using a Simple Mail Transfer Protocol (SMTP) Server, ensure that you configure the SMTP Server before you start a restore operation. Set the configuration in Microsoft Management Console (MMC) by right-clicking **Dashboard** in the navigation tree and selecting **Properties**. Then, in the **E-mail** property page, enter the SMTP server and port.

- **Exchange Server 2013:** Install the correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects on the Exchange server from which you are running the mailbox restore operations.

Download and install the Exchange MAPI and Microsoft Outlook MAPI on different servers. Do not install Microsoft Outlook 2010 or 2013 on the same server that Data Protection for Microsoft Exchange Server uses for mailbox restore operations. Conflicts might occur in the MAPI configurations.

- **Exchange Server 2016 or later:** Install Microsoft 32-bit Outlook 2016 or later versions as the MAPI client on the same server that Data Protection for Microsoft Exchange Server uses for mailbox restore operations.

About this task

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a .pst file. When you restore a mailbox to the mailbox restore destination, IBM Spectrum Protect Snapshot automatically restores the mail items in the Recoverable Items folder.
 - You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the mail items in the folders.
 - The mail items that you can restore depends on whether the mailbox is enabled for mailbox restore operations.
 - You can restore the Recoverable Items content for a public folder mailbox but not for each public folder in the public folder mailbox.
 - You can exclude the mail items in the Recoverable Items folder in mailbox restore operations.
 - You cannot create a subfolder in the Recoverable Items folder in a mailbox.
 - The **Mailbox Restore Browser** displays folders that are normally hidden from view, for example, in the Recoverable Items folder. Folder names in the Recoverable Items folder are internal to Microsoft Exchange and are not translated by Microsoft. Therefore, if you use a language other than English, the folder names still display in English.
- In Exchange Server 2016 or later, when opening a mailbox in **Mailbox Restore Browser** view, the mailbox needs to be restored to a temporary mailbox first. The amount of time that it takes to complete the restore process depends on the size of the mailbox databases, and the network speed. Do not open multiple mailboxes at the same time to avoid long delays.
- In Exchange Server 2013, you can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder. However, you cannot restore individual messages in a public folder by using the Mailbox Restore Browser interface.
 - To restore an Exchange 2013 public folder mailbox, the Exchange user must have the Public Folders management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange Server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange Server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange Server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.

If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.
 - You might restore to a different public folder mailbox than the original mailbox if, for example, the public folder is relocated after the time of the backup. Before you complete the public folder restore operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.
- If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all the mailboxes are in the same recovery database.
- By default, IBM Spectrum Protect Snapshot restores the latest backup that is available for the specified mailbox.

Restriction: Only mailboxes within the same database can be restored in a single mailbox restore action.

Procedure

1. Start MMC.

2. Under the **Protect and Recover Data** node in the navigation tree, select **Exchange Server**.
3. On the **Recover** panel, click **View > Mailbox Restore Browser**.
4. In the **Select Source** window, specify the mailbox that you want to restore.

Restriction: A list of all available user mailboxes in the domain is displayed, including those mailboxes that were not backed up. Mailboxes that are not backed up cannot be selected for a restore operation. Only mailboxes that are backed up can be restored.

Choose from the actions in the following table:

<i>Table 21. Selecting mailboxes to restore</i>	
Task	Action
Browse mailboxes and select one to restore	<ol style="list-style-type: none"> a. From the drop-down list, select Browse Mailboxes. b. Select a mailbox. c. Click OK. <p>Tip: Use the Search field to filter the mailboxes. You can also sort the mailboxes by columns.</p>
Specify a mailbox to restore by name	<ol style="list-style-type: none"> a. In the Mailbox Name field, enter the name of the mailbox to restore. b. Click OK.
Restore a mailbox backup that was created at a specific time	<ol style="list-style-type: none"> a. In the Backup Date/Time field, click the default date and time to edit the details. b. To change the date, click the calendar icon, select a date, and press Enter. c. To change the time of day, use the 12-hour system convention such as 2 p.m. d. Click OK.
Review the mailbox backups that are available to restore before you complete the restore operation	<ol style="list-style-type: none"> a. From the drop-down list, select Browse Mailboxes. b. Select a mailbox for which backups exist. c. From the Available Database Backups list, review the backups that are available for the mailbox and select a backup version to restore. d. Ensure that the Backup Date/Time field reflects the time stamp for the selected mailbox backup. e. Click OK.
Restore a mailbox that was deleted or re-created after the time of the backup	<p>In the Actions pane, click Properties, and on the General page, enter the temporary mailbox alias.</p> <p>Tip: If you do not enter the alias, the mailbox restore operation uses the administrator mailbox as a temporary storage location.</p>

<i>Table 21. Selecting mailboxes to restore (continued)</i>	
Task	Action
Browse all databases in a backup	a. From the drop-down list, select Browse Databases . b. Select a database. c. Click OK . Tip: Use the Search field to filter the databases. You can also sort the mailboxes by columns.

After the selected mailbox is restored to the recovery database, the restored mailbox and folders are displayed in the results pane.

5. In the results pane, browse the folders and messages that are contained within the selected mailbox. Choose from the following actions to select the mailbox, folder, or message to restore:

<i>Table 22. Previewing and filtering mailbox items</i>	
Task	Action
Preview mailbox items	a. Select a mailbox item to display its contents in the preview pane. b. When an item contains an attachment, click the attachment icon to preview its contents. Click Open or save the item by clicking Save .
Filter mailbox items	Use the filter options to narrow the list of folders and messages in the result pane. a. Click Show Filter Options and Add Row . b. Click the down arrow in the Column Name field and select an item to filter. You can filter by folder name, subject text, and so on. You can filter public mailbox folders only by the Folder Name column. When you select All Content , the mailbox items are filtered by attachment name, sender, subject, and message body. c. In the Operator field, select an operator. d. In the Value field, specify a filter value. e. If you want to filter on more items, click Add Row . f. Click Apply Filter to filter the messages and folders.

6. In the **Actions** pane, click the folder or messages restore task that you want to run.

If you click **Save Mail Message Content**, which becomes available only when a message is selected in the preview pane, a Windows Save File window is displayed. Specify the location and message name and click **Save**.

The **Restore Progress** window opens and shows the progress of the restore operation. IBM Spectrum Protect Snapshot restores the mailbox backup to its original mailbox location.

7. To restore a mailbox or mailbox item to either of the following locations, complete the following steps. Choose from the actions in the following table:

<i>Table 23. Restoring a mailbox to another mailbox or .pst file</i>	
Task	Action
Restore a mailbox or mailbox item to a different mailbox	<p>a. On the Actions pane, click Open Exchange Mailbox.</p> <p>b. Enter the alias of the mailbox to identify it as the restore destination.</p> <p>c. Drag the source mailbox to the destination mailbox on the results pane.</p> <p>Restriction: You cannot drag mail items or subfolders in the Recoverable Items folder to a destination mailbox.</p>
Restore a mailbox to an Outlook personal folders (.pst) file	<p>a. On the Actions pane, click Open non-Unicode PST File (for Exchange Server 2013) or Open Unicode PST File (for Exchange Server 2016 or later).</p> <p>b. When the Windows File window opens, select an existing .pst file or create a .pst file.</p> <p>c. Drag the source mailbox to the destination .pst file on the results pane.</p>
Restore Public Folder Mailbox	<p>Select this action to restore a public folder mailbox to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore. If you are restoring a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\) at the end of the folder path.</p> <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>

In the Actions pane, the **Close Exchange Mailbox** and **Close PST File** tasks are displayed only when a destination mailbox or .pst file is opened.

8. Optional: Remove the recovery database by clicking **Close Mailbox to Restore**.

This option is displayed only after a recovery database is created. IBM Spectrum Protect Snapshot removes the recovery database and cleans up the restored files. If you do not select **Close Mailbox to Restore**, the recovery database is not removed even if you exit MMC.

If MMC also detects a recovery database that is created outside of IBM Spectrum Protect Snapshot, it automatically connects to it. When you complete your mailbox restore tasks, you must manually remove the recovery database. You cannot use the **Close Mailbox to Restore** option.

Related concepts

Exchange mailbox restore operations

By using IBM Spectrum Protect Snapshot, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

Restoring mailboxes directly from Exchange Server database files

If the backup database (EDB) file and log files are available on the disk of a supported Microsoft Exchange Server, you can restore an individual mailbox directly from the EDB file.

Before you begin

When you restore mailboxes directly from Exchange Server database files, verify that you have read and write access to the EDB file, and verify that the Exchange Server transaction log files exist.

If you use IBM Spectrum Protect for Virtual Environments software, review the following guidelines before you restore the mailbox:

- You can use IBM Spectrum Protect for Virtual Environments to back up an Exchange Server in a virtual machine. For more information about the **backup** command, see [Backup command \(http://www.ibm.com/support/knowledgecenter/SSERB6_8.1.4/ve.user/r_ve_vmcli_backup.html\)](http://www.ibm.com/support/knowledgecenter/SSERB6_8.1.4/ve.user/r_ve_vmcli_backup.html).
- To restore mailboxes from the backups that are created by IBM Spectrum Protect for Virtual Environments, mount the virtual volumes that contain the EDB file and log files with read/write access. You can obtain read/write access by clearing the **Mount virtual volume as read only** check box.
- If the log files are on a different volume than the EDB file, mount the volume that contains the log files on an unused drive letter. In this way, you can apply the transaction logs to the EDB file.

Procedure

1. From the Exchange Server, start IBM Spectrum Protect Snapshot.
2. After you log on to IBM Spectrum Protect Snapshot, in the navigation area, select the **Exchange Server** node and the **Recover** tab.
The **Mailbox Restore Browser** view opens.
3. In the **Actions** pane, click **Open EDB File on Disk**.
4. In the window, enter or browse to the location of the EDB file.
5. In the window, enter or browse to the location of the log file directory.
Specifying a path to the log file directory is not required. However, the amount of time that is necessary to complete the restore operation is reduced when you provide the log file directory path.
6. Click **OK**.
The EDB file is opened and the mailboxes are displayed.
7. Select the mailbox that you want to restore and the type of restore that you want to complete.
For example, you can restore a mailbox to a PST file.
8. When the restore operation is complete, click **Close Mailbox to Restore**.
You are prompted to save or delete the recovery database folder.

Restoring a deleted mailbox or items from a deleted mailbox

You can use IBM Spectrum Protect Snapshot for Microsoft Exchange Server to restore a mailbox or mailbox items that were deleted from an Exchange Server.

Before you begin

Decide where the mailbox data from the deleted mailbox is to be restored.

If you restore mail to a Unicode personal folder (.pst) file, or you restore a mailbox that is deleted or re-created after the time of the backup, IBM Spectrum Protect Snapshot for Microsoft Exchange Server

requires a temporary mailbox to store the mailbox messages. Create a temporary mailbox by setting the Alias of temporary mailbox option on the **Properties** page, under the **General** tab.



Attention: Ensure that the temporary mailbox that you create is on a database with enough disk storage capacity to accommodate all of the mailbox items that you are restoring.

Procedure

Complete one of the following actions:

- Restore the deleted mailbox data to the original location. Before you run the mailbox restore operation, re-create the mailbox that is using Exchange.
- Restore the deleted mailbox data into an active alternative mailbox in an online Exchange Server.
- Restore the deleted mailbox data into an Exchange Server personal folders (.pst) file.

Restoring mailboxes on remote systems

The process of restoring mailboxes on a remote system with the **Mailbox Restore Browser** feature differs from restore operations on local systems.

Before you begin

- For a typical mailbox restore task, you must install the IBM Spectrum Protect Snapshot for Microsoft Exchange Server package on the local and remote systems. The correct version of Microsoft Exchange Server MAPI Client and Collaboration Data Objects (for Exchange Server 2013) or 32-bit Outlook (for Exchange Server 2016 and later) must also be installed.
- To restore mailboxes on a remote system with **Mailbox Restore Browser**, the local and remote systems must be in the same domain. The following procedure assumes that you installed the latest version of IBM Spectrum Protect Snapshot, configured the Exchange Server workload, and have a Windows PowerShell remote connection. At least one mailbox needs to be stored in a database on the remote system.
- After you complete the installation of software on the local and remote systems, verify that the remote system is available and that you can connect to it. Verify that the database with the mailbox you want to restore is backed up successfully. You can use Microsoft Management Console (MMC) to go to the remote system where you want to restore mailboxes.

Procedure

1. From MMC, expand the navigation tree to the remote system.
2. From the **Protect and Recover Data** tree node, select the **Exchange Server**.
3. In the main window, on the **Recover** tab, click **View > Mailbox Restore Browser**.
4. Select the mailbox that you want to restore. Click **OK**.
The mailbox is displayed in the **Source Mailbox** tree view.
5. Click **Restore Mail to Original Mailbox**.

Protecting SQL Server data

With IBM Spectrum Protect Snapshot for SQL Server, you can back up and restore Microsoft SQL Server databases in a stand-alone configuration.

About this task

You can create point-in-time snapshots of your Microsoft SQL Server and store the data locally on the server that is running the backup.

Note: For information about protecting Microsoft SQL Server data in VMware environments, see chapter *Protection for in-guest applications* in the *IBM Spectrum Protect for Virtual Environments: Data Protection for VMware User's guide*.

Related concepts

Prerequisites

With IBM Spectrum Protect Snapshot, you can back up and restore SQL Server data and protect your SQL Server environment.

Related tasks

Configuring

You can use configuration wizards to configure IBM Spectrum Protect Snapshot, or you can complete the steps manually. For best results, be guided by the step-by-step instructions in the configuration wizards.

Prerequisites

With IBM Spectrum Protect Snapshot, you can back up and restore SQL Server data and protect your SQL Server environment.

Security requirements for backup and restore operations

IBM Spectrum Protect Snapshot for SQL Server requires certain settings to process backup and restore operations in a secure environment.

To install IBM Spectrum Protect Snapshot for SQL Server, you must have Windows administrator authority. You must register IBM Spectrum Protect Snapshot for SQL Server to the IBM Spectrum Protect server and you must use the appropriate node name and password when it connects to the IBM Spectrum Protect server.

You can specify SQL Server logon information in one of the following ways:

- Accept the default sa account and system administrator password. Ensure that you secure your sa login account with a password.
- Use SQL user ID security and specify both the SQL user name and password. With SQL user ID security, the SQL Server administrator provides the logon ID and the password that provides access to the SQL Server.
- Use a trusted connection and allow Windows authenticate the logon.

You must add the SQL logon user name or Windows user name to the SQL Server SYSADMIN fixed server role before IBM Spectrum Protect Snapshot for SQL Server can use those credentials.

Choosing your Microsoft SQL Server backup strategy

Depending on your SQL Server environment, you can run full backups only, copy-only full backups, full plus log backups, full plus differential backups, or file and group backups. Your backup strategy might also be to back up data to IBM Spectrum Protect or local shadow volumes.

Full backup method (Legacy and VSS)

Use the full backup method for system databases such as *master*, *model*, and *msdb* because of their typical small size. A full backup can take a long time to run. However, the restore process is the most efficient because only the most recent full backup is restored.

Copy-only full backup method (Legacy and VSS)

Use the copy-only full backup method to periodically create copy-only full backups for long-term retention without affecting existing backup schedules or retention policies that you use for disaster recovery. Copy-only full backups do not affect the transaction logs or the sequence of backups, such as differential backups or full backups.

Full backup plus log backup method (Legacy and VSS)

Use the full backup plus log backup method when the normal backup schedule or network capacity cannot support a full backup.

To minimize the effect on the backup schedule and network traffic during peak times, you can run a periodic full backup, followed by a series of log backups. For example, you can schedule full backups on

the weekend and log backups during the week. You can run full backups during low usage times and when increased network traffic can be tolerated.

Restriction: If you run multiple full backups, the SQL database log can become full. Subsequent backups might fail as a result. If necessary, use basic SQL Server tools to truncate the log of your SQL databases.

Full backup plus differential backup method (Legacy and VSS)

Use the full backup plus differential backup method if your backup schedule and network capacity can facilitate backing up all transaction logs that accumulate between full backup operations. This strategy requires that only one differential backup plus the last full backup be transferred to complete a restore operation. However, the same amount of data must be transferred in the differential image, as in the series of log backup operations.

Although you can run only VSS full backups, you can apply legacy differential backups to the VSS full backup.

Full backup plus differential plus log backup method (Legacy and VSS)

Use the full backup plus differential plus log backup method to reduce the number of transactions that must be restored and applied. Restore operations are faster as a result.

If, for example, you complete a full legacy or VSS backup weekly, and a differential backup nightly, and a log backup every four hours, the restore processing would include the full backup, a differential backup, and at most five log backups. However, if you only complete a full plus log backup scheme on the same cycle, the restore processing would include a full backup plus up to 41 log backups (six days multiplied by six log backups per day plus up to five backups on the day the full backup is completed).

Although you can run only VSS full backups, you can apply legacy log backups and legacy differential backups to the VSS full backup.

File or group backup method (only Legacy)

Use the file or group backup method when it is impractical to back up an entire database because of the size of the data, or associated time and performance issues.

When a group is created on the SQL Server, database files are identified with that group. The group that is used for the group backup depends on the group to which the database files are defined.

File or group options can save backup and restore processing time when certain tables or indexes have more updates than others and must be backed up more often. It is time-effective to contain such data in their own file group or files, and to back up only those items.

Except for logical log files, you can back up your transaction logs after you back up a data file or file group.

IBM Spectrum Protect backups versus local shadow volumes backups

When you create a policy for your backups, you must choose whether to back up data to IBM Spectrum Protect storage versus VSS disks. Data backups to IBM Spectrum Protect typically takes longer to process than backups to local shadow volumes.

Backing up SQL Server data to IBM Spectrum Protect is necessary when long-term storage is required. For example, saving SQL Server data on tape for archival purposes requires long-term storage. IBM Spectrum Protect backups are also necessary for disaster recovery situations when the disks that are used for local backups are unavailable.

By maintaining multiple backup copies on IBM Spectrum Protect server storage, a point-in-time copy is available if backups on the local shadow volumes become corrupted or deleted.

Restriction: If you run legacy log backups to an IBM Spectrum Protect server, the SQL database log files can be truncated.

Local shadow volumes

When you back up data to local shadow volumes, ensure that sufficient local storage space is assigned to the local shadow volumes. Create different sets of policies for backups to both local shadow volumes and to IBM Spectrum Protect server storage. If you use a VSS provider other than the Windows VSS System Provider, follow the backup recommendations of the VSS provider.

You can run backups to local shadow volumes by time and backup versions. It is more effective to base policy for local backups on version limits because local snapshots are created more frequently and VSS storage provisioning and space limitations apply. In AlwaysOn Availability Group (AAG) environments, all of the AAG members must use the same local VSS policy.

Environment and storage resources also impact how many backup versions you can maintain on local shadow volumes for VSS fast restore and VSS instant restore operations, and on IBM Spectrum Protect server for VSS restore operations.

Data protection for Microsoft SQL AlwaysOn Availability Groups

You can run VSS (full) and legacy (full, differential, file/set/group, and log) backup operations on a primary replica. You can run copy-only VSS and legacy backup operations, and normal log backups on a secondary replica. You cannot run a differential backup on a secondary replica.

For backups on a secondary replica, the replica must be in the synchronized or synchronizing state. You can have multiple AlwaysOn Availability Groups (AAGs) in an SQL Server cluster. You can also have a mix of standard databases and AAGs on an SQL Server cluster.

When you back up data, you can distribute the backup workload for scalability and isolate backup activity to a dedicated backup node. When you isolate backup activity, it minimizes the effect on production databases.

Given that replicas are copies of the same database, avoid redundant backups of the same databases. Apply retention policies to unique databases.

As a best practice, allow backups from any node in the availability group and enable restore operations from any node in the availability group.

Best practices for backing up data in an AAG

When you use IBM Spectrum Protect Snapshot for SQL Server to manage AAG backups, consider the following backup options:

Backup priority

Specified per database in an AAG, the backup priority option defines the order in which replicas are used to back up a database in an AAG.

Preferred replica

Specified at an AAG level, the preferred replica option defines whether primary or secondary replicas can be used for backup operations.

- Prefer secondary replica: Scheduled backups occur on a secondary replica, if available. If the secondary replica is not available, you can use the primary replica.
- Secondary only: Scheduled backups can occur only on a secondary replica.
- Primary: Scheduled backups can occur only on the primary replica.
- Any replica: Scheduled backups can occur on any replica.

/USEALWAYSONNode parameter

A parameter option on the **backup** command that provides a common namespace for all backups. Each node authenticates separately with IBM Spectrum Protect. Backed up data is stored in the AlwaysOnNode namespace by using the Asnode option.

/ALWAYSONPriority parameter

A parameter option on the **backup** command that specifies that a local availability database is backed up only if it has the highest backup priority among the availability replicas that are working properly. This parameter applies only to scheduled backups.

Typical data protection deployments in AAG environments

You can back up data in an AAG in the following ways:

- Distribute a legacy backup across AAG replicas
- Distribute a VSS backup across AAG replicas

Approach®: Legacy backups are distributed across AAG replicas

When you configure your environment to distribute a legacy backup across AAG replicas, follow these steps:

1. Set the preferred replica to **Prefer secondary replica**.
2. Install IBM Spectrum Protect Snapshot for SQL Server on all replicas that are eligible to run a backup.
3. Create a command script to run a .CMD file with a **backup** command similar to the following example:

```
tdpsqlc backup db1,db2,db3 full /alwaysonpriority
```

4. Associate each IBM Spectrum Protect Snapshot for SQL Server node with the defined schedule.
5. Run backups on the SQL node according to defined priorities for each database.

Scenario: VSS backups are distributed across AAG replicas

When you configure your environment to distribute a VSS backup across AAG replicas, follow these steps:

1. Set the preferred replica to **Prefer secondary replica**.
2. Install IBM Spectrum Protect Snapshot for SQL Server on all replicas that are eligible to run a backup.
3. Create a command script to run a .CMD file with a separate **backup** command per database similar to the following sample

```
tdpsqlc backup db1 full /alwaysonpriority /backupmethod=VSS
      backupdest=TSM
tdpsqlc backup db2 full /alwaysonpriority /backupmethod=VSS
      backupdest=TSM
tdpsqlc backup db3 full /alwaysonpriority /backupmethod=VSS
      backupdest=TSM
```

4. Associate each IBM Spectrum Protect Snapshot for SQL Server node with the defined schedule.
5. Run backups on the SQL node according to defined priorities for each database.

Preparing for VSS instant restore operations

By using the VSS instant restore feature, you can restore one or more databases from a VSS snapshot backup on local shadow volumes that are managed by IBM Spectrum Protect Snapshot. Data is restored at the volume level.

About this task

In a VSS instant restore operation, you can restore full backup and copy backup types. For SQL, custom application, and file system data, you can run VSS instant restore operations in a Microsoft Failover Clustering environment. You cannot use parallel VSS restore operations on Microsoft Windows Server.

Procedure

1. Verify that your system has a VSS hardware provider with VSS instant restore capability, for example, IBM XIV VSS Hardware Provider or System Storage support for Microsoft Volume Shadow Copy Service software.
2. Verify that you are restoring local VSS backups of SAN-attached volumes on the same storage systems where the backups are located.

You can run instant restore operations on the following storage systems:

- IBM System Storage DS8000 series
 - IBM System Storage SAN Volume Controller
 - IBM Storwize
 - IBM XIV Storage Systems
 - IBM and non-IBM storage devices that implement the Microsoft VSS ResyncLuns API
3. Verify that databases are restored to the same drive letter and paths that are used during the original backup.
 4. Close applications or windows that might have files or handles open on the volumes that are being restored.

Verifying the integrity of legacy databases by using the checksum option

With IBM Spectrum Protect Snapshot, you can verify the integrity of legacy database backups by setting a checksum option.

About this task

A *checksum* is a value that is calculated and written in the data page header of the database data file. When a data file is read again, the checksum value is recalculated. Checksum processing validates the values in a file or configuration for unexpected changes. Values are verified between the current state and the baseline state.

Restriction: Checksum integrity checking is available only with legacy backups on SQL Server.

Procedure

1. Open the **General Properties** window in Microsoft Management Console (MMC).
2. Select **Compute SQL Server checksum for legacy backup**.

If you select this option, all legacy backups are checked by default. You can override this setting to set integrity checking for a particular backup.

For example, if you bypassed integrity checking on all backups, you can set integrating checking on a particular legacy backup by selecting the **SQL Checksum** backup option on the **Protect** tab for the SQL instance. You can also issue the SQLCHECKSUM option with the **backup** command on the command line to temporarily enable or disable the checksum option.

Results

When you select the **Compute SQL Server checksum for legacy backup** check box, the setting is written to the IBM Spectrum Protect Snapshot for SQL Server preferences file, `tdpsql.cfg`, and is applied to all legacy backup operations. If you clear the check box, integrity checking does not apply to any legacy database backup.

Related tasks

[“Creating legacy backups of SQL Server databases” on page 125](#)

You can create a legacy backup of your standard SQL databases by using Microsoft Management Console (MMC). You can also use the legacy method to back up availability databases with SQL Server 2012 and later versions.

Creating VSS backups of SQL Server databases

You can back up standard SQL Server databases or availability databases by using Microsoft Volume Shadow Copy Service (VSS).

Before you begin

- Before you back up a database, run **dbcc checkdb** and **dbcc checkcatalog** to verify the logical and physical consistency of the database.

- To manage local VSS backups or to run offloaded backups to IBM Spectrum Protect server storage, ensure that IBM Spectrum Protect Snapshot is configured in your environment.

If you use VSS to back up data to an IBM Spectrum Protect server, IBM Spectrum Protect Snapshot is not required.

- Do not mix persistent and nonpersistent VSS snapshots.

About this task

On SQL Server 2012 and later versions, you can back up availability databases in an AlwaysOn Availability Group (AAG) regardless of which availability replica is used for the backup operation.

Restriction: When you complete a full backup of a secondary replica in an AAG, only a copyfull backup of that database is created.

To back up availability databases, ensure that IBM Spectrum Protect Snapshot is configured to use an AlwaysOn node. Additionally, specify the AlwaysOn node in the **AlwaysOn Node** field in the **TSM Node Names** page of the IBM Spectrum Protect Configuration Wizard. If you change the **AlwaysOn node name** field in the **AlwaysOn Node** properties page for your SQL workload, you must run the IBM Spectrum Protect Configuration Wizard to complete the reconfiguration of the name.

If you do not want to use the IBM Spectrum Protect Configuration Wizard to register the node on the IBM Spectrum Protect server, you can use the IBM Spectrum Protect **register node** command.

Restriction: You cannot back up a temporary database because a temporary database is created each time the SQL Server starts.

Procedure

1. Start Microsoft Management Console (MMC).
2. If you plan to use offloaded backups, and your environment is configured for use with an IBM Spectrum Protect server, specify a value in the **Remote DSMAGENT Node name** field.
 - a) Select the **SQL Server** instance in the navigation tree, and click **Properties** in the **Actions** pane.
 - b) Select the **VSS Backup** property page. If the **Remote DSMAGENT Node name** is blank, enter a node name.

An offloaded backup uses another system (specified with the **Remote DSMAGENT Node name** parameter) to move SQL data to IBM Spectrum Protect server storage. Offloaded backups can reduce the load on network, I/O, and processor resources during backup processing.

3. On the **Protect** tab of an SQL instance, select an option for viewing databases.

<i>Table 24. Database backup views</i>	
Task	Action
View a list of SQL databases that are available for a backup operation	Click View: Databases .
View a list of SQL Server 2012 and later version availability databases that are available for a backup operation	Click Standard Databases . Information about the availability databases in an availability group is displayed, including the replica role, synchronization state, and space and log usage. Toggle the Standard Databases / Availability Databases button for the respective database views.

Refine the list of available databases in the results pane by entering a keyword in the **Search** field. Then, select the databases to back up.

4. Verify the backup options. If the backup options are not displayed, click **Show Backup Options**.
If you want to use offloaded backups, select **Yes** in the **Offload** field.

5. In the **Actions** pane, click **Backup Method** and select **VSS**.
6. In the **Actions** pane, click **Backup Destination** and select a location to store the backup:

Local

Click this item to store the database backups to only local shadow volumes.

TSM

Click this item to store the database backups only on IBM Spectrum Protect server storage. Do not select this option if you are using IBM Spectrum Protect Snapshot in a stand-alone configuration.

Both

Click this item to store the database backups to IBM Spectrum Protect server storage and local shadow volumes. Do not select this option if you do not have an IBM Spectrum Protect Snapshot license or if you are using IBM Spectrum Protect Snapshot in a stand-alone configuration.

If you set the **Backup Destination** to TSM, a nonpersistent VSS snapshot is created. To avoid mixing persistent and nonpersistent VSS backups, do not follow a series of backups to a local server with a backup to IBM Spectrum Protect server. As a best practice, set the **Backup Destination** to BOTH to send data to IBM Spectrum Protect server and preserve the local snapshot backup versions.

7. Optional: Choose a mode for the current task:
 - **Run Interactively:** Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
8. To start the backup operation, in the **Actions** pane, take one of the following actions:
 - a) Click **Full Backup**.
Alternatively, right-click a database and select the backup action that you require from the menu.
 - b) Click **Copy-Only Full Backup**.
A copy-only full backup is independent of the sequence of SQL Server backups, and is not used as a base for a differential backup. A differential backup is not associated with the copy-full backup, but is associated with the previous full backup that completed. You might use a copy-only full backup as a special purpose backup that does not affect backup and restore operations, and retain such a backup for longer than conventional backups.
9. Review the status of the backup operation by clicking **Task List** in the results pane. Click **Task Details** to view detailed status information.

Results

During backup processing, IBM Spectrum Protect Snapshot for Microsoft SQL Server bypasses database snapshots and databases that are in offline, mirroring, and restoring states.

What to do next

To determine which databases backups are bypassed during backup processing, review the `tdpsql.log` in the directory where IBM Spectrum Protect Snapshot is installed.

Creating legacy backups of SQL Server databases

You can create a legacy backup of your standard SQL databases by using Microsoft Management Console (MMC). You can also use the legacy method to back up availability databases with SQL Server 2012 and later versions.

Before you begin

- For legacy database backups, you can verify whether a backup is valid without physically restoring the backup. Before you restore the legacy database backup, you can run the restore operation with the **Verify Only** option in Microsoft Management Console (MMC).
- To run a legacy backup, ensure that the IBM Spectrum Protect Snapshot for SQL Server license file is installed.

- On SQL Server 2012 and later versions, you can also back up availability databases in an AlwaysOn Availability Group (AAG) regardless of which availability replica is used for the backup operation. To back up availability databases, ensure that IBM Spectrum Protect Snapshot is configured to use an AlwaysOn node. Additionally, specify the AlwaysOn node in the **AlwaysOn Node** field in the **TSM Node Names** page of the IBM Spectrum Protect Configuration Wizard.

Procedure

1. Start MMC.
2. Select the **SQL Server** instance in the tree view.
3. On the **Protect** tab of an SQL instance, select an option for viewing databases.

<i>Table 25. Database backup views</i>	
Task	Action
View a list of SQL databases that are available for a backup operation	Click View: Databases .
View a list of SQL Server 2012 and later version availability databases that are available for a backup operation	Click Standard Databases . Information about the availability databases in an availability group is displayed, including the replica role, synchronization state, and space and log usage. Toggle the Standard Databases / Availability Databases button for the respective database views.

Refine the list of available databases in the results pane by entering a keyword in the **Search** field.

4. Verify the backup options. If the backup options are not displayed, click **Show Backup Options**.

<i>Table 26. Database backup options</i>	
Option	Action
Data Stripes	Use this option to specify the number of data stripes to use in a backup or restore operation. The <i>numstripes</i> variable can be in the range 1 - 64. The default value is 1. When you use a multiple stripes number for legacy backups, and set the Verify Only parameter to Yes to restore the legacy backup, the number of stripes for the legacy restore must be equal to or greater than the number of stripes for the legacy backup.

Table 26. Database backup options (continued)	
Option	Action
Estimated Database % Change	<p>Use this option to specify the estimated percentage of the database that changed since its last full database backup. The default value is 20.</p> <p>This estimate is useful because SQL Server does not provide a way to determine the size of a differential backup, and because the IBM Spectrum Protect server requires an accurate size estimate to efficiently allocate space and place objects. The IBM Spectrum Protect server uses this value to determine whether there is enough space in the primary storage pool to contain the backup.</p>
Estimated Log % Change	<p>Use this option to specify the estimated percentage of an SQL database that changed due to non-logged operations since the last log backup. The default value is 0.</p>
Truncate Logs	<p>Use this option to specify whether to dispose of entries that you no longer need in the SQL database transaction log after you back up the log. The default value is Yes.</p> <p>In general, you do not want to truncate the log when you rebuild a corrupted database. This option enables the server to back up the transaction log but does not affect the data. All transaction log entries are written from the time of the last log backup to the point of database corruption. If you do not truncate the transaction log, you might be able to back up the transaction log of a damaged, suspect, or unrecoverable SQL Server database.</p>
Back Up Tail-Log	<p>Use this option to store log records that are not backed up.</p> <p>By storing these records, also known as the <i>tail of the log</i>, the log chain is kept intact. Before you can recover an SQL Server database to the last point in time, you must back up the tail of the transaction log. The tail-log backup is the last backup of interest for the database recovery plan.</p>

Table 26. Database backup options (continued)	
Option	Action
SQL Server Checksum	<p>Use this option to verify the integrity of a legacy database backup. Integrity checking is a process that validates the values in a file or configuration for unexpected changes. Values are verified between the current state and the baseline state.</p> <p>In the Performance Properties window of MMC, you can enable or disable the checksum option for all your legacy databases at once. You can override the global setting, and temporarily enable or disable the checksum option for a database backup, by setting this SQL Checksum option to Yes or No.</p>

5. In the **Actions** pane, click **Backup Method** and select **Legacy**.
6. Optional: Choose a mode for the current task:
 - **Run Interactively**: Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled**: Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
7. To start the backup operation, in the **Actions** pane, take one of the following actions:
 - **Full Backup**
 - **Copy-Only Full Backup**
 - **Differential Backup to TSM**
 - **Log Backup to TSM**
8. Review the status of the backup operation by clicking **Task List** in the results pane. Click **Task Details** to view detailed status information.

What to do next

- To determine which databases backups are bypassed during backup processing, review the `tdpsql.log` in the directory where IBM Spectrum Protect Snapshot is installed. IBM Spectrum Protect Snapshot bypasses database snapshots and databases that are in offline, mirroring, and restoring states.
- To determine whether the checksum option is applied to a legacy database backup, enter the **tdpsqlc query tsm *** command on the command line, or the equivalent **Get-DpSqlBackup** cmdlet.

Related tasks

[“Verifying the integrity of legacy databases by using the checksum option” on page 123](#)

With IBM Spectrum Protect Snapshot, you can verify the integrity of legacy database backups by setting a checksum option.

Cloning an SQL Server database

Clone standard SQL Server production databases or availability databases by using Microsoft Volume Shadow Copy Service (VSS) and Windows PowerShell cmdlets. You can create a persistent VSS snapshot of the production database at any time. This snapshot is a clone of the production database. Use this clone to complete database activities you would normally do on the production database without impacting the production database. Use Windows PowerShell cmdlets to query the cloned database, and to mount and unmount it on remote servers.

Before you begin

- You must install and enable Windows PowerShell 3.0, or later, on all IBM Spectrum Protect Snapshot installations that you are using. To download, install, and enable the software, follow the instructions in: [Microsoft Windows Management Framework 3.0 Downloads](#).
- The database that you are cloning must be on a storage device that is managed by a VSS hardware provider.
- Before you create a clone, ensure that sufficient space is available on the storage device.

About this task

- You can create as many clones of the production database as you want, providing sufficient storage space is available on your storage device. Each clone is a snapshot on your storage device.
- You can create a clone on the same SQL Server where the active database is running only.
- You can mount the clone to remote servers so that other users can access and use the clone. Clones are always mounted as read/write.
- You can access and use the clone on any SQL Server that is attached to the storage device.
- You must delete a clone manually as it does not expire based on policy settings.

Procedure

1. Open a Windows PowerShell command prompt.
2. Create a clone of your active production database by issuing the following Windows PowerShell cmdlet. You can specify the **name** parameter to specify the name of your production database. When created, you can access and work with all the files that are contained in the cloned database without any impact to the production database.

```
New-CloneFromDpSqlComponent
```

Tip: At any time, you can query the cloned databases that are created on the current SQL server and node by issuing the following Windows PowerShell cmdlet. You can view all details about all the clones that are created. You can also specify a particular database name to query all the clones that are created for that database.

```
Get-DpSqlClone
```

3. Mount the created clone to any server (that is running IBM Spectrum Protect Snapshot Version 8.1.4) by issuing the following Windows PowerShell cmdlet. You can mount the clone as often as you want to one or more servers, but only to one at a time. You can also mount the clone to remote servers that other users can access. The cloned database is always mounted as read/write.

```
Mount-DpSqlClone
```

You can use the **-PostProcessScript** parameter to specify the name of a command script that you can use to run operations on cloned databases after they are successfully mounted; for example, to attach the mounted database to an SQL Server instance. For the script name, you can specify either a relative directory path (starting from your current working directory) or the full directory path. You can also provide additional parameters for the specified command script. However, only the directory path name of the command script can contain backslash characters. If the command script fails, the cloned database is still mounted, but a warning message is displayed. See the following examples:

```
Mount-DpSqlClone -name db1
-PostProcessScript .\mypostscript.ps1
```

```
Mount-DpSqlClone -name db1
-PostProcessScript c:\path\mypostscript.ps1
```

```
Mount-DpSqlClone -name db1 -PostProcessScript
"c:\path\mypostscript.ps1 -parm1 val1 -parm2 val2"
```

4. Unmount the clone from a server by issuing the following cmdlet.

```
Dismount-DpSqlClone
```

You can use the **-PreProcessScript** parameter to specify the name of a command script that you can use to run operations on mounted cloned databases before dismounting; for example, to detach a mounted database from an SQL Server instance. For the script name, you can specify either a relative path (starting from your current working directory) or the full directory path. You can also provide additional parameters for the specified command script. However, only the path name of the command script can contain backslash characters. If the command script fails, the unmount operation terminates. See the following examples:

```
Dismount-DpSqlClone -mountpoints  
c:\mnt1 -PreProcessScript .\myprescript.ps1
```

```
Dismount-DpSqlClone -mountpoints  
c:\mnt1 -PreProcessScript c:\path\myprescript.ps1
```

```
Dismount-DpSqlClone -mountpoints  
c:\mnt1 -PreProcessScript "c:\path\myprescript.ps1 -parm1 val1 -parm2 val2"
```

5. Unlike backups, cloned databases do not expire based on policy settings and remain on your system until they are deleted manually. When you no longer require the cloned database, delete it from the SQL Server by issuing the following cmdlet.

```
Remove-DpSqlClone
```

Restriction: The delete operation is possible only from the system where the clone was created.

Tip: To view the details about a specific cmdlet, issue the **Get-Help** cmdlet with the cmdlet name. To enhance how the information is displayed (particularly in languages other than English), display the help in a separate window by including the **-showwindow** parameter with the help command.

```
Get-Help Get-DpSqlClone -showwindow
```

Related tasks

[“Enabling Windows PowerShell Remoting for Remote Management and Remote Mounting” on page 155](#)

From a single IBM Spectrum Protect Snapshot installation, you can manage all of the IBM Spectrum Protect Snapshot installations in your organization.

Related reference

[“Cmdlets for protecting Microsoft SQL Server data” on page 162](#)

The following table identifies the cmdlets that you can use to protect Microsoft SQL Server data.

Deleting SQL Server backups

You can remove an SQL Server backup that you created with the VSS backup method. Complete this task only if necessary.

Before you begin

Typically, backups are deleted automatically based on user-defined policy management settings. This procedure is necessary only if you must delete backups that are outside the scope of IBM Spectrum Protect Snapshot policy definitions.

If you delete a remotely-mounted backup, the snapshots and the relationship between the source and target volumes on the storage device are also deleted. However, the target volume that is imported and mounted might continue to exist. In addition, the target volume might not be available to the server where the remote mount occurred. The operations to the target volume depend on the VSS hardware provider and the storage device implementation.

After the maximum number of remotely-mounted backup versions or the maximum number of days to retain a backup is exceeded, the associated backup is expired and deleted.

Procedure

1. Start Microsoft Management Console (MMC).
2. Click **Recover Data > SQL** in the **Management** window.
3. On the **Recover** tab for the SQL instance, select **View: Database Restore**. In the results pane, browse to and select one or more database backups to delete.

The corresponding node type, for example, DP or AlwaysOn, must also be selected.

4. In the **Actions** pane, click **Delete backup**.

When a backup is deleted, two tasks display in the task window to show you that the deletion is in progress, and that the view is being refreshed.

Related tasks

[Mounting VSS snapshots to remote servers](#)

You can mount VSS snapshots to remote servers that other users can access. When you remotely mount a snapshot, you must use hardware that includes a feature for creating transportable snapshots. If you use a mix of hardware, all hardware must include a feature for creating transportable snapshots.

Deactivating legacy backups of SQL Server databases

IBM Spectrum Protect deactivates an SQL database backup as a part of IBM Spectrum Protect policy management. Data backups are typically deactivated when an SQL database is deleted from the SQL Server as part of the scheduled backup processing.

Before you begin

The SQL database that you want to deactivate must be a legacy backup. You cannot use this procedure to deactivate VSS backups. The **Delete** action is available in the **Actions** pane when you select a VSS backup from the **Recover** view.

About this task

For legacy backups, you can deactivate any or all of the following backup object types: full, differential, copyfull, log, file, group, or set. You can also deactivate any object or object type that is older than a specified number of days.

When you deactivate database backups, all copy group parameters, which control how backup versions are generated, located, and expired, are inspected. Any existing backups on IBM Spectrum Protect server are subject to deletion, as specified by relevant policy settings.

Tip: After a full SQL backup, all preceding copy-only full, file, group and differential backups stop adhering to the **VERExists** and **RETEExtra** settings, even if the databases still exist on the Data Protection for SQL Server client system. In the management class for these backup objects, set the **VERDeleted** and **VERExists** parameters to the same value and also set the **RETEExtra** and **RETOOnly** parameters to the same value to maintain consistent version-expiration behavior.

When automatic processing is insufficient, the **inactivate** function explicitly deactivates one or more active data backups on the IBM Spectrum Protect server.

Procedure

1. Under the **Protect and Recover Data** node in the tree view, select the SQL Server.
2. Open the **Recover** view to see the status of the backup.
Active backups are displayed.
3. Select the database backup that you want to deactivate, and in the **Actions** pane, click **Inactivate**.
4. To view the results, take one of the following actions:
 - To display the database that you made inactive, click **All Backups** on the toolbar.
 - To display only active database backups, click **Active Backups** on the toolbar.

Setting single-user mode for restore operations

You might have to start an SQL Server instance in single-user mode during certain restore operations. For example, you might use single-user mode when you are restoring a damaged master database or a system database, or when you are changing server configuration options.

Before you begin

Restriction:

- You cannot restore SQL databases that are in use. By placing SQL databases to be restored in single-user mode, you can avoid system attempts to restore those databases.
- Microsoft Management Console (MMC) cannot connect to a SQL Server instance that is started in single-user mode. If you want to use MMC when the SQL Server instance is in single-user mode, you must use the command-line interface, `tdpsqlc.exe`, to restore the master database.

Procedure

1. To determine which users are using the databases, use the SQL stored procedure, `SP_WHO`.
2. To force users off the SQL database and set the SQL Server to single-user mode, issue this `TRANSACT-SQL` command.

```
ALTER DATABASE DBNAME SET SINGLE_USER
```

```
WITH ROLLBACK AFTER N SECONDS
```

3. To start the SQL Server in single-user mode, use the `-m SQL SERVER` startup option.
4. To return the database to multiple-user mode, issue this `TRANSACT-SQL` command.

```
ALTER DATABASE DBNAME SET MULTI_USER
```

Setting data restore options

To optimize the data restore process for your environment, modify the default options that are available in Microsoft Management Console (MMC).

Procedure

1. On the **Recover** tab, select **Database Restore**.
2. Click **Show Restore Options** to modify the default restore options as follows:

Option	Action
Auto Select	Specify a value of Yes (default) to select the backup objects to restore. With automatic selection, when you select the most recent backup to restore, all associated backups are automatically selected, up to the previous full backup. This option affects backups in the following ways: When you click a differential backup, the associated full backup is also selected. When you click a log backup, the associated full backup and all associated earlier differential or log backups are also selected.
Performance	

<i>Table 27. Database restore options. (continued)</i>	
Option	Action
Stripes	<p>Specify the number of data stripes to use in a restore operation. A maximum of 64 data stripes is allowed. The default value is 1.</p> <p>The value that you enter must correspond to the value that you set for SQL buffers.</p> <p>Restriction: This restore option is available only with legacy backups.</p> <p>When you use multiple stripes for legacy backups, with the Verify Only parameter set to Yes, the number of stripes for legacy restore must be equal to or greater than the number of stripes for that backup.</p>
Restore Behavior	
Database Owner Only	<p>To mark a database for owner use only, set this value to Yes. The default value is No, specifying not to mark the database for owner use.</p> <p>Restriction: This restore option is only available with legacy backups.</p>
Keep CDC	<p>For databases enabled for change data capture (CDC), set this value to Yes to retain, during a legacy restore operation, the change data capture records that with recorded changes. These changes are insertions, deletions, and edits to SQL Server database tables. The default value is No.</p> <p>Restriction: This restore option is only available with legacy backups and applies to all legacy backup types except for log backups.</p>
Replace	<p>To replace a database during a restore operation, set this value to Yes. The default value is No, which specifies not to replace databases.</p> <p>Restriction: This restore option is available only with legacy backups.</p>
Recovery	<p>Use this option to restore data to an SQL database that is not on a standby SQL Server. The default value is Yes.</p> <ul style="list-style-type: none"> • Select Yes when you run a sequence of restore operations to an SQL database and the current restore operation is the final one in the sequence. Alternatively, use this option when it is the only restore operation. • Select No when you run a sequence of restore operations to an SQL database and the current restore operation is not the final one in the sequence. Select No for all restore operations in the sequence except for the final one.

<i>Table 27. Database restore options. (continued)</i>	
Option	Action
Stand By Undo File Name	<p>For this option, specify a value of Yes to change the target SQL database to a standby SQL database. The default value is No.</p> <p>This option is available for full, differential, and log backup types. When you specify this option for a database, it applies to all backup objects for that database. Similarly, when you remove this option for a backup object, the option is removed for all backup objects.</p>
Verify Only	<p>Before you restore a legacy database backup, set this option to Yes to verify that the database backup volumes are complete and that all can be read. The default value is No.</p> <p>Note: This option verifies only that the database backup volumes are complete and that they are readable. It does not verify the structure of the data that is contained in the backup volumes.</p> <p>When you use multiple stripes for legacy backups, with this option set to Yes for legacy backups, the number of stripes for the restore must be equal to or greater than the number of stripes for that backup.</p> <p>Restriction: This restore option is available only for legacy database backups.</p>
Source Server	
From SQL Server	<p>Use this option to specify the name of the SQL Server that the backup is created from.</p> <p>To specify the name of a virtual environment SQL Server, change IncludeTsmVM to Yes to view Virtual Environment backup SQL databases in the Databases view. The backup method is listed as TSMVM to distinguish these databases from the other databases that are listed.</p>
Tape	
Wait for Tape Mounts for Restore	<p>Use this option to specify whether the IBM Spectrum Protect Snapshot for Exchange Server restore operation waits for the IBM Spectrum Protect server to mount removable media such as tapes or other sequential device media. The default value is Yes.</p>
Wait for Tape Mounts for File Information	<p>When you query IBM Spectrum Protect for file information, use this option to specify if IBM Spectrum Protect Snapshot for Exchange Server waits for the IBM Spectrum Protect server to mount removable media. The default value is Yes.</p> <p>Restriction: You can use this restore option only with legacy backups.</p>
VSS	

<i>Table 27. Database restore options. (continued)</i>	
Option	Action
Instant Restore	<p>Specify a value of Yes to use volume-level snapshot restore (instant restore) for local VSS backups if the backup exists on SAN-attached volumes. Specify a value of No to disable instant restore, which bypasses volume-level copy and uses file-level copy (fast restore) to restore the files from a local VSS backup. The default value is Yes, which uses volume-level snapshot restore if it is available.</p> <p>This option is available for VSS operations only. Instant restore for SAN Volume Controller earlier than Version 5.1 or DS8000, you must ensure that previous background copies that involve volumes that are being restored are completed first.</p> <p>In an instant restore operation, files on the destination file system are overwritten. Incremental and differential backups are automatically converted to file-level restores. An instant restore operation requires that the drive or volume where the mailbox database is located must be available. Any other process or application must not have access to the drive or volume.</p>

Related tasks

“Troubleshooting VSS backup and restore operations” on page 174

If you encounter a problem during VSS backup and restore processing, attempt to reproduce the problem in your environment.

Restoring SQL Server data

You can restore SQL Server databases or parts of databases only from `full`, `copyfull`, `differential`, and `log` backups. You can also restore availability databases with SQL Server 2012 and later versions.

Before you begin

If multiple instances of SQL Server are running, ensure that you specify the server name in IBM Spectrum Protect Snapshot for SQL Server to access the correct SQL Server.

About this task

Restriction: You cannot restore VSS backups to an alternate SQL Server. When you restore a database, existing data is overwritten by the restored data and is no longer available after the restore operation is complete.

- The Regional settings, which are defined in the **Regional** property page, must match the date format that is defined for the Microsoft SQL Server.
- You can use VSS to run backup operations of type `full` or `copyfull`. You can apply legacy differential and legacy log backups after a full VSS backup is restored.
 - When Virtual Environment restore operations are configured from the IBM Spectrum Protect server, you can restore and view these databases from the Recover tab.
 - You can also restore availability databases that you backed up with the `AlwaysOn` node with SQL Server 2012 and later versions. Backups of availability databases can be restored to any availability replica in an availability group.
 - You can restore a legacy database backup that is verified as valid and complete with the **Verify Only** option in Microsoft Management Console (MMC).

Procedure

1. Start MMC.
2. Select the **SQL Server** instance in the tree.
3. On the **Recover** tab for the SQL instance, specify the type of SQL data to restore.

<i>Table 28. Database backup views</i>	
Task	Action
View a list of SQL databases that are available for a restore operation	Click View: Databases .
View a list of SQL database backup files that are available for a restore operation	Click View: Files .
View a list of SQL Server 2012 and later version availability databases that are available for a restore operation	Click DP Node Backups to show AlwaysOn node backups. Toggle the DP Node Backups / AlwaysOn Node Backups button for the respective database views.

4. On the **Recover** tab of an SQL Server instance, select an option for viewing databases.
In the **Results** pane, browse to the databases that are available to restore. The following options are available:

<i>Table 29. Database restore selection options</i>	
Option	Action
Search	Enter a keyword in the Search field to refine and filter the list of databases.
Filter	Use the filter options to refine and filter the list of databases. <ol style="list-style-type: none"> a. Click Show Filter Options and Add Row. b. In the Column Name field, click the down arrow and select an item to filter. c. In the Operator field, select an operator. d. In the Value field, specify a filter value. e. If you want to filter on more items, click Add Row. f. Click Apply Filter.
Backups	Select the database to restore. You can click Active Backups to show only active backups, or click All Backups to show both active and inactive backups.
Refresh	Click Refresh to update the view with your changes.

If you applied a filter, the objects on the server that match the filter or search criteria are listed on the **Recover** tab. The status area indicates the number of items that match the criteria n of x displayed, where n equals the number of objects that match the filter criteria, and x is the number of objects that are retrieved from the server. For example, 5 of 20 displayed. If you specify refresh options to further narrow your results, and click **Refresh** again, the objects on the server that match the filtered and refresh options are displayed. Each time that you click **Refresh**, another query is run against the IBM Spectrum Protect server.

5. Verify the options for the restore operation. If the restore options are not displayed, click **Show Restore Options**.
6. Optional: Choose a mode for the current task:
 - **Run Interactively:** Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
7. To start the restore operation, in the **Actions** pane, take one of the following actions:
 - Click **Restore**.
 - Click **Restore VerifyOnly**. The **Restore VerifyOnly** task is available only if all the selected database backups are legacy backups.

Important: When you select the **Restore VerifyOnly** action, the number of stripes that are used for the restore must be the same or greater than the number of stripes that are used for the backup you are verifying. If it is not, the **Restore VerifyOnly** operation terminates with an error.
8. To view the status of the restore operation, click **Task List** in the results pane. Click **Task Details** to view detailed status information.

Related reference

Regional

Use the **Regional** property page to set preferences that affect how languages and information are displayed and logged.

Restoring an SQL Server database to an alternate instance

Using Microsoft Management Console (MMC), you can restore an SQL Server database backup to an alternate SQL Server instance or database. You can also restore availability databases to an alternate location on any availability replica in an availability group. Restore to an alternate instance by using Microsoft Management Console (MMC), Windows PowerShell cmdlets, or the command-line interface (CLI).

Before you begin

Install IBM Spectrum Protect Snapshot for SQL Server on both systems.

About this task

You can also restore availability databases that you backed up with the AlwaysOn node. Backups of availability databases can be restored to any availability replica in an availability group.

You can select only one database at a time when you restore a database to an alternate location.

Procedure

1. Copy the IBM Spectrum Protect Snapshot for SQL Server options file (`dsm.opt`) from the source system to the target system.

Source system

The system from which the original backup (to be restored) is created.

Target system

The alternate system to which the backup is to be restored.

By default, the `dsm.opt` file is in the `C:\Program Files\Tivoli\TSM\TDPSql` directory. If you specified a value of `generate` for the **passwordaccess** parameter in the `dsm.opt` file, you might need to reset the password for this node on the IBM Spectrum Protect server.

2. Start MMC.
3. On the **Recover** tab for the SQL instance, specify the type of SQL data to restore.

<i>Table 30. Database backup views</i>	
Task	Action
View a list of SQL databases that are available for a restore operation	Click All Backups .
View a list of SQL Server 2012 and later version availability databases that are available for a restore operation	Click DP Node Backups to show AlwaysOn node backups. Toggle the DP Node Backups / AlwaysOn Node Backups button for the respective database views.

4. Verify the options for the restore operation. If the restore options are not displayed, click **Show Restore Options**.
 - a) Ensure that **Wait for Tape Mounts for Restore** is set to **Yes**.
 - b) Ensure that **Wait for Tape Mounts for File Information** is set to **Yes**.
 - c) If the database to be restored is to replace an existing database on the target system, click **Replace**.
 - d) Use the **Instant Restore** option to turn Instant Restore on or off. Click **Yes** to use Instant Restore. Click **No** to disable Instant Restore if you want to use Fast Restore.



Attention: Instant Restore operations overwrite all files on the destination file system.

5. To start the backup operation, in the **Actions** pane, take one of the following actions:

- a) Click **Restore to Alternate Location**.
- b) Click **Restore VerifyOnly to Alternate Location**.

The **Restore VerifyOnly to Alternate Location** task is available only if all the selected database backups are legacy backups.

Important: When you select the **Restore VerifyOnly to Alternate Location** action, the stripes number must be the same or greater than that which is set in the backup. If it is not, the **Restore VerifyOnly** operation terminates with an error.

6. In the **Restore Into** section of the **Alternate Location Restore Settings** window, click **Restore to new database**, and specify a target SQL Server instance name and target database name to restore a backup object to.

The **Database name** can have the same name as the source database or you can specify a different unique name.

For VSS backups, the only instance available to restore to is that which you select on the **Recover** tab before starting the backup operation.

7. In the **Relocate** section of the window, filter the restore processing operations.

<i>Table 31. Database backup views</i>	
Task	Action
Specify new destination locations in which to restore backed up SQL databases, logs, and FILESTREAM files (SQL Server 2008 or later versions)	Click Restore all files into one directory .
Restore the log files into a location that is different from where the SQL database and other related files are restored	Select Relocate logs into and specify a new path in the text entry field.

<i>Table 31. Database backup views (continued)</i>	
Task	Action
Restore FILESTREAM files (SQL Server 2008 or later versions) into a location that is different from where the SQL database and logs are restored relevant for legacy restore operations only	Select Relocate other files into , and specify a new path in the text entry field.
Restore one or more individual SQL database, log, and FILESTREAM files relevant for legacy restore operations only	Click Relocate files individually , and click Browse to open a folder selection window. Select a folder or create a new folder, and click OK . The path of the selected files entries is set to use the folder. This option is available for legacy backups only.

Restriction: You cannot relocate database files and logs with a partial restore operation in MMC. You must use the command-line interface to complete a partial restore operation that requires these parameters.

8. Click **Restore** to close the **Alternate Location Restore Settings** window and begin the restore.
9. To view the status of the restore operation, click **Task List** in the lower half of the results pane. Click **Task Details** to view detailed status information.

What to do next

You can restore a legacy database backup that is verified as valid and complete with the **Verify Only** option in MMC, or with the **/VERIFYOnly** option of the **restore** command on the command line.

Restoring the master database

A damaged master database can prevent the SQL Server from starting and cause other errors. To protect your data if the master database is damaged, you must routinely complete a full database backup of the master database (msdb).

Before you begin

- Set single-user mode for restore operations.
- Always keep an up-to-date backup of your master database because the master database contains the system catalog. The system catalog contains important information about the SQL Server configuration.
- Ensure that you back up the master database after any changes that update system tables. For example, back up the master database after you use any of these statements:
 - ALTER DATABASE
 - CREATE DATABASE
 - DISK INIT
 - DISK RESIZE
 - DISK MIRROR
 - DISK UNMIRROR
 - DISK REMIRROR
 - Various DBCC options such as SHRINKDB
 - System-stored procedure, such as sp_droptremotelogin, sp_addumpdevice, sp_dropdevice, sp_addlogin, sp_droplogin, sp_addserver, sp_dropserver, sp_addremotelogin

About this task

You must complete a VSS restore of the master database while the database is offline. Therefore, you must stop the associated SQL Server instance before you run the restore operation. If you restore an online master database, the operation might fail or disable subsequent VSS backup and VSS restore operations until the SQL Server VSS Writer service is restarted.

If the master database is damaged while a server instance is running, fix the damaged database by restoring a recent full master database backup. If a server instance cannot start because the master database is damaged, the master database must be rebuilt. When you rebuild a master database, all system databases revert to their original state.

Restriction: Microsoft Management Console cannot connect to an SQL Server instance that is started in single-user mode. When the SQL Server instance is in single-user mode, you must use the command-line interface, `tdpsqlc.exe` to restore the master database.

Procedure

1. Click **Start > All Programs > IBM Spectrum Protect > Data Protection for Microsoft SQL Server > SQL Client - Command Line**.
2. Start the SQL Server in single-user mode.
3. Use IBM Spectrum Protect Snapshot for SQL Server to restore the master database.
When the master database finishes the restoration process, the SQL Server shuts down and an error message is displayed. The message indicates that the connection to the SQL Server is lost. This loss of connection is expected.
4. Restart the database engine to restore SQL Server to the typical multiuser mode.
5. Run the SQL Server setup program to rebuild the master database.
When you rebuild the master database, use the same character set and sort order as the master database backup that is to be restored.
6. Manually reapply any changes to the master database that occurred after the date of the database backup that is used to complete the restore operation.
7. Restore the msdb database.
During the process of rebuilding the master database, the SQL Server setup program drops, and then re-creates, the msdb database. Therefore, you must restore the msdb database with the master database.

Results

After the master database is restored, you can use MMC to back up and restore individual databases that are operating in single-user mode.

Related tasks

[Troubleshooting VSS offline restore of a master database](#)

Microsoft SQL Server only supports offline VSS restores of the master database. IBM Spectrum Protect Snapshot for SQL Server does not support offline restore operations. Therefore, you cannot use IBM Spectrum Protect Snapshot for SQL Server to restore the master database.

[“Setting single-user mode for restore operations” on page 132](#)

You might have to start an SQL Server instance in single-user mode during certain restore operations. For example, you might use single-user mode when you are restoring a damaged master database or a system database, or when you are changing server configuration options.

Restoring SQL databases with full-text catalogs and indexes

You can restore SQL Server 2005 and 2008 databases, including their full-text catalogs and full-text indexes.

About this task

When you back up an SQL Server 2005 database and the full-text index is part of a full-text catalog, the full-text catalog has a physical path. In this scenario, the full-text catalog is treated as a database file.

When you back up an SQL Server 2008 database and later data, a full-text catalog is either a logical or virtual object that contains a group of full-text indexes. This full-text catalog does not have a physical path. When you restore a database with SQL Server 2008 and later full-text catalog files, no data is explicitly stored. The file is automatically backed up and restored as part of the filegroup.

Procedure

- To restore a database with the SQL Server 2005 physical full-text catalog file from the command-line interface, use the **/RELocate** and **/TO** parameters.

For example:

```
Restore DATABASE full /relocate=database,sysft_docindex,database_log
```

```
/TO={database_dir}\database.mdf,{database_dir}\docindex,
```

```
{database_log_dir}\database_log.ldf
```

- To restore a database with the SQL Server 2005 physical full-text catalog file from the GUI, use the **Relocate files individually** option.

From the command-line interface, use **/relocate** and **/TO** instead of **/RELOCATEDir**.

Protecting SQL Server data in a Windows Server Core environment

Server Core is a minimal and low-maintenance server environment where you can run the minimum services that are necessary to maintain Windows Server 2008 and later versions. You can install and operate IBM Spectrum Protect Snapshot in this minimal server environment.

About this task

You can install and use IBM Spectrum Protect Snapshot on Windows Server 2008 R2 Server Core SP1 and later versions.

In such a minimal environment, only the command-line interface is available for IBM Spectrum Protect Snapshot on Windows Server Core unless you use the IBM Spectrum Protect Snapshot remote management support. Additionally, if you use Windows Installer (MSI) to install IBM Spectrum Protect Snapshot, you can use only the unattended mode.

You can use the **backup** and **restore** commands to protect databases that are stored on Microsoft SQL Server 2012 or later versions.

Backing up SQL Server databases on Windows Server Core

To back up Microsoft SQL Server databases, use the **backup** command.

About this task

Use the following procedure to back up SQL Server databases to the IBM Spectrum Protect server, or to take local VSS snapshots.

Procedure

1. To back up all or part of an SQL database on Windows Server Core, enter the following command at the command prompt:

```
tdpsqlc backup database_name backup_type [other_options]
```

where *database_name* specifies the name of the database, and *backup_type* specifies the type of backup such as a full backup. You can specify other options, such as the back up method.

For example, to create a full legacy backup of SQL databases DB_01 and DB_02, enter the following command:

```
tdpsqlc backup DB_01,DB_02 full /backupmethod=legacy
```

For example, to create a full legacy backup of all databases on the SQL Server, enter the following command:

```
tdpsqlc backup * full /backupmethod=legacy
```

2. To back up a file group, enter the following command at the command prompt:

```
tdpsqlc backup database_name file_group
```

where *database_name* specifies the name of the database, and *file_group* specifies the file group in the database.

For example, to back up the filegroup DB_01_group1 that belongs to the DB_01 database, enter the following command:

```
tdpsqlc backup DB_01 Group=DB_01_group1
```

Restoring SQL Server databases on Windows Server Core

To restore Microsoft SQL Server databases, use the **restore** command.

About this task

Use the following procedure to recover all or part of one or more SQL databases.

Procedure

- To restore all or part of an SQL database on Windows Server Core, enter the following command at the command prompt:

```
tdpsqlc restore database_name backup_type [other_options]
```

where *database_name* specifies the name of the database, and *backup_type* specifies the type of backup such as a full backup. You can specify other options, such as the file group.

For example, to create a full database restore of databases DB_01 and DB_02, and to replace the existing databases with the database objects that are recovered from the IBM Spectrum Protect server, enter the following command

```
tdpsqlc restore DB_01 group=DB_01_group1
```

To restore the filegroup DB_01_group1 that belongs to the DB_01 database, enter the following command:

```
tdpsqlc restore DB_01 group=DB_01_group1
```

To restore all the logical files that are in the DB_01 database, enter the following command:

```
tdpsqlc R DB_01 file=*
```

Changing IBM Spectrum Protect Snapshot configuration values on Windows Server Core

To configure preferences for IBM Spectrum Protect Snapshot for SQL Server, use the **set** command at the Windows Server Core command prompt.

About this task

The values that you change are saved in the IBM Spectrum Protect Snapshot configuration file. The default configuration file is `tdpsql.cfg`.

Procedure

- At the command prompt, enter the following command:

```
tdpsqlc set parameter=value [/configfile=filename]
```

where *parameter* is the IBM Spectrum Protect Snapshot parameter or option for which you want to change the value, and *value* is the new value that you want to specify. **/configfile** is the optional parameter for the configuration file name. If you do not specify the **/configfile** parameter, the default configuration file (`tdpsql.cfg`) is used.

Examples:

Task

Set the preferred SQL Server in the `tdpsql.cfg` file.

Command: `tdpsqlc set sqlserver=your_SQL_instance /configfile=tdpsql.cfg`

Command: `tdpsqlc set fromsqlserver=your_SQL_instance /configfile=tdpsql.cfg`

Task

Change the name of the IBM Spectrum Protect Snapshot activity log file to `tdpsql.log`.

Command: `tdpsqlc set logfile=tdpsql.log`

Protecting custom application and file system data

With IBM Spectrum Protect Snapshot, you can back up and restore custom application and file system data, and protect your environment.

About this task

When you use IBM Spectrum Protect Snapshot software that is configured with a IBM Spectrum Protect server, and you create VSS snapshot backups of application and file system data, you can send the data to the IBM Spectrum Protect server storage pools. The data is set as a full image backup. The backup that is stored on the IBM Spectrum Protect server is used to restore volumes and mount points.

Regularly back up the `x:\adsm.sys\vss_staging` directory, the `LSM_REP_LOG_VOL_SNAP`, and `.TsmFmDatabases` folders. If the folders become damaged and unusable, restore the last copy. When the `TargetSetsState` file has a mapping entry, the reconciliation process can detect an inconsistency between `TargetSetsState` and the metadata that is stored on the IBM Spectrum Protect server, and synchronize the data. For example, if there is an orphan `VSSDC_XX` value in the IBM Spectrum Protect server and the value does not exist in `TargetSetsState`, reconciliation processing deletes the orphan metadata from the IBM Spectrum Protect server.

Related tasks

[Configuring IBM Spectrum Protect Snapshot for file system and custom applications in a Microsoft Cluster Server environment](#)

You can configure IBM Spectrum Protect Snapshot to support a file system and custom application workload in a Microsoft Cluster Server environment.

Prerequisites

To create VSS snapshot backups of NTFS or ReFS file systems and applications, use IBM Spectrum Protect Snapshot. When you back up applications and file systems, IBM Spectrum Protect Snapshot must access the data.

If permissions must be granted for IBM Spectrum Protect Snapshot to access the data, see the documentation that is provided with the application and file system.

Scripts for automated processing

You can run scripts to prepare and resume custom application and file systems before and after you create a snapshot of the data.

To prepare custom application and file systems for volume-level snapshots, you can use preprocessing (**PRESNAPSHOTCMD**) and postprocessing (**POSTSNAPSHOTCMD**) scripts. If specified, these scripts run during backup processing. For example, you can use the **PRESNAPSHOTCMD** script to quiesce an application and the **POSTSNAPSHOTCMD** to resume it.

Data protection in an environment with IBM Spectrum Protect server

IBM Spectrum Protect Snapshot provides a way for you to manage persistent snapshots on Windows file systems by using VSS backup operations.

You can use IBM Spectrum Protect to protect custom application data and file systems in the following ways:

- Restore VSS snapshots of file systems when the backup destination is set IBM Spectrum Protect server (TSM option), a local system (LOCAL option), or both destinations (BOTH option). For backups to TSM or BOTH destinations, a VSS snapshot of the selected drive letter or mount point is created. The backup is sent to the IBM Spectrum Protect server as an image-level backup of the VSS snapshot.
- Restore backups that are in IBM Spectrum Protect server pools at a file system or drive level, as an image-level restore of the VSS snapshot.
- Restore backups that are in IBM Spectrum Protect server storage to an alternate server location, different local drive letter, or mount point.
- For backups to TSM or BOTH destinations, complete offloaded backups from a secondary system to IBM Spectrum Protect server storage pools.
- From the command-line interface and Microsoft Management Console (MMC), use enhanced final backup summary statistics for information about client-side deduplication, compression, encryption, and other options.
- From the command-line interface and Microsoft Management Console (MMC), use enhanced query output for information about client-side deduplication, compression, and encryption usage.
- For backups to TSM or BOTH destinations, use the unified views of available file system backup versions.

VSS backups are managed as backup versions by IBM Spectrum Protect Snapshot management policies. VSS backups remain available for VSS Instant Restore or VSS Fast Restore operations. When IBM Spectrum Protect server is available in the environment, you can use the IBM Spectrum Protect backup-archive client to create file-level backups of your file system or custom application data in IBM Spectrum Protect storage pools.

To import VSS snapshots, verify that the VSS provider can use transportable snapshots. To use the command-line interface for the mount command with remote options, verify that the VSS provider can use transportable snapshots and configure the Windows PowerShell Remoting feature.

Backing up custom application and file system data

You can back up custom application and file system data by using Microsoft Volume Shadow Copy Service (VSS).

Before you begin

Configure IBM Spectrum Protect Snapshot to manage VSS snapshots for the custom application or file system. by using the **Standalone Configuration Wizard**. In the wizard, select **File System**.

Alternatively, if you are using IBM Spectrum Protect server, configure integration with the IBM Spectrum Protect server by using the IBM Spectrum Protect wizard.

Restriction: If you configure IBM Spectrum Protect Snapshot to integrate with IBM Spectrum Protect server, do not simultaneously configure the items listed.

- In the VSS Requestor options file (baclient\dsm.opt), do not specify the following entry:

```
Include.Image volume management-class-name
```

- In the IBM Spectrum Protect Snapshot configuration file (fcmcfg.xml), Exchange configuration file (tdpexc.cfg, or SQL configuration file (tdpsql.cfg), do not specify VSSPOLICY statements that use the TSM option to back up data to IBM Spectrum Protect.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the tree view, click **Protect and Recover Data > File System**.
3. In the **Protect** tab, select the volume names and mount points to back up.

Tip: Enter a keyword in the **Search** field to refine the list of available volume names and mount points in the results pane.

4. Click **Show Backup Options**. Then, select one of the options and take the appropriate action.

Option	Action	More information
Optional: For custom applications, specify the presnapshot and postsnapshot batch scripts	Specify the complete path for a presnapshotcmd file or postsnapshotcmd file to use. These scripts are used to quiesce or stop the application, which is necessary to ensure backup consistency.	A presnapshotcmd file is a Windows command file that is run before a snapshot backup is created. For example, the presnapshotcmd script can quiesce an application before the snapshot is created. A postsnapshotcmd file is a Windows command file that is run after a snapshot backup is created. For example, a postsnapshotcmd script can resume an application after the snapshot is created.

Table 32. Backup options (continued)		
Option	Action	More information
To use offloaded backups	Set the Offload option. If you intend to use offloaded backups, ensure that the Remote DSMAGENT Node name field is complete when you set preferences for the Data Protection properties. If you use the command-line interface to update the configuration for offloaded backups, set the REMOTEDSMAGENTNODE parameter. This parameter applies only to VSS backups.	An offloaded backup uses another system to move custom application and file system data to IBM Spectrum Protect server storage. An offloaded backup can reduce the load on the network, I/O, and CPU resources during backup processing.

5. Optional: Choose a mode for the current task:
 - **Run Interactively:** Click this item to run the current task interactively. This selection is the default.
 - **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.
6. In the **Actions** pane, click **Backup Method** and select **VSS**.
7. In the **Actions** pane, click **Backup Destination** and specify where to store the backup.
These options are available:
 - **Local**
 - **TSM**
 - **Both**
8. In the **Actions** pane, click **Full Backup**. You can view the backup progress from the **Task List** and **Task Details** panes.

Implementing custom application and file system backup scenarios

You might decide to implement different backup strategies that depend on your network traffic requirement and backup schedule. Follow these typical backup scenarios if you want to back up custom application and file system data to local shadow volumes, or create a file-level backup to the IBM Spectrum Protect server.

Creating a VSS snapshot backup to local shadow volumes

You can back up custom application and file system data to local shadow volumes if sufficient storage space is available on the local shadow volumes.

Procedure

1. If you are backing up an application, stop (or suspend) the custom application and file system.
To complete this task automatically, use the **fcmlcli backup** command and specify the **/presnapshotcmd= cmdstring** parameter where the *cmdstring* variable is the command that runs run before the snapshot operation begins.
2. Create the backup by specifying the **fcmlcli backup** command. Or, in the IBM Spectrum Protect Snapshot user interface, click **Full Backup to Local** option in the **Action** window.

VSS snapshot backups can also be sent to the IBM Spectrum Protect server storage pools by specifying **BackupDestination TSM**, or BOTH from either the command-line interface (with the **fcmlcli backup** command) or IBM Spectrum Protect Snapshot user interface. The VSS snapshot data is sent to the IBM Spectrum Protect server as an image-level backup.

3. If you are backing up an application, restart or resume the application.

To complete this task automatically, use the **fcmlcli backup** command and specify the **/postsnapshotcmd= cmdstring** parameter where the *cmdstring* variable is the command that runs run after the snapshot operation ends.

Related concepts

[Choosing your Microsoft SQL Server backup strategy](#)

Depending on your SQL Server environment, you can run full backups only, copy-only full backups, full plus log backups, full plus differential backups, or file and group backups. Your backup strategy might also be to back up data to IBM Spectrum Protect or local shadow volumes.

Related reference

[Backup examples](#)

These output examples provide a sample of the text, messages, and process status that displays when you use the **backup** command.

[Backup optional parameters](#)

Optional parameters follow the **backup** command and positional parameters.

Creating a file-level backup to the IBM Spectrum Protect server

You can optionally use the IBM Spectrum Protect server backup-archive client to create file-level backups of your file system or custom application data.

Procedure

1. Create a VSS snapshot backup.
2. Mount the VSS snapshot backup.
3. Create a file-level backup to the IBM Spectrum Protect server by issuing an IBM Spectrum Protect backup-archive client command. Use the incremental, selective, or archive command with the snapshotroot option. The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data that is created by a volume snapshot.
4. Unmount the VSS snapshot backup.

Related tasks

[Mounting custom application and file system backups](#)

You can mount a snapshot backup to see a point-in-time consistent copy of custom application and file system data.

[Creating a VSS snapshot backup to local shadow volumes](#)

You can back up custom application and file system data to local shadow volumes if sufficient storage space is available on the local shadow volumes.

Mounting custom application and file system backups

You can mount a snapshot backup to see a point-in-time consistent copy of custom application and file system data.

About this task

Restriction: You cannot use Microsoft Management Console (MMC) to mount a backup to a different server. To mount a VSS snapshot to a remote server, enter the **mount backup** command at the command prompt.

When you submit a mount request, all of the volumes that are contained in the original snapshot set are imported. If the number of volumes that are imported exceed the maximum number of allowable mapped volumes for the environment, the mount operation can fail.

Procedure

1. Start MMC.

2. Click **Recover Data** in the welcome page of MMC.
3. In the **Recover** tab, go to the **Action** pane. Click **Mount Backup**.
4. Either type the path to the empty NTFS or ReFS folder where you want to mount the backup or browse to find the path. Click **OK**.
On the **Recover** tab, the backup that you mounted is displayed.
5. Use the **Explore** and **Unmount Backup** options in the **Actions** pane to complete tasks with the backup that you mounted.

Deleting custom application and file system backups

You can remove a custom application or file system VSS backup object that you created with the VSS backup method. Complete this task only if necessary.

Before you begin

Typically, backups are deleted automatically based on user-defined policy management settings. This procedure is necessary only if you need to delete backups that are outside the scope of your standard policy management definitions.

Procedure

1. Start Microsoft Management Console (MMC).
2. From the **Management** window, click **Protect and Recover Data > File System**.
3. On the **Recover** tab, select the volume name or mount point to delete.
Be aware that you are not deleting the volume or mount point. You are deleting the backup version of the volume or mount point. To view active and inactive backups, click **All Backups**. To view only active backups, click **Active Backups**.
4. Right-click to select the volume or mount point; then, either click **Delete Backup** in the menu, or click **Delete Backup** in the **Actions** pane.
A confirmation message is displayed.
 - To delete the volume, click **Yes**.
 - To stop the deletion process, click **No**.

When a backup is deleted, two tasks are displayed in the task window to show that the deletion is in progress, and that the view is being refreshed.

Restoring custom application and file system data

The IBM Spectrum Protect Snapshot user interface displays information about active and inactive backups. Review this information so that you can select the custom application and file system data to restore.

About this task

When you submit a restore request, all of the volumes that are contained in the original snapshot set are imported. If the number of volumes that are imported exceed the maximum number of allowable mapped volumes for the environment, the restore operation can fail.

Procedure

1. Start Microsoft Management Console (MMC).
2. In the **Management** window, click **Protect and Recover Data > File System**.
3. On the **Recover** tab, select an option for viewing databases.
In the **Results** pane, browse to the databases that are available to restore. The following options are available:

<i>Table 33. Database restore selection options</i>	
Option	Action
Search	Enter a keyword in the Search field to refine and filter the list of databases.
Filter	Use the filter options to refine and filter the list of databases. <ul style="list-style-type: none"> a. Click Show Filter Options and Add Row. b. In the Column Name field, click the down arrow and select an item to filter. c. In the Operator field, select an operator. d. In the Value field, specify a filter value. e. If you want to filter on more items, click Add Row. f. Click Apply Filter.
Backups	Select the database to restore. You can click Active Backups to show only active backups, or click All Backups to show both active and inactive backups.
Refresh	Click Refresh to update the view with your changes.

If you applied a filter, the objects on the server that match the filter or search criteria are listed on the **Recover** tab. The status area indicates the number of items that match the criteria n of x displayed, where n equals the number of objects that match the filter criteria, and x is the number of objects that are retrieved from the server. For example, "5 of 20 displayed." If you specify refresh options to further narrow your results, and click **Refresh** again, the objects on the server that match the filtered and refresh options are displayed. Each time that you click **Refresh**, another query is run against the IBM Spectrum Protect server.

4. Verify the restore options. If the restore options are not displayed, click **Show Restore Options**.

Set a value for the following options:

FromServer

If the backup is not displayed in the results pane, enter the name of the server where the original backup was completed. The default value is the current server.

InstantRestore

To use VSS Instant Restore, enter Yes. This option applies only to snapshots that are on a disk system that supports Instant Restore operations. Enter No to use VSS Fast Restore (file-level copy).

5. Optional: Choose a mode for the current task:

- **Run Interactively:** Click this item to run the current task interactively. This selection is the default.
- **Run Scheduled:** Click this item to convert the current action into a scheduled task. When you select this item, the schedule wizard runs the command that is required to complete the task.

6. In the **Actions** pane, complete one of the following actions to begin the restore operation.

- a) Click **Restore** to restore the selected volumes or mount points to their original location.
- b) Click **Restore Into**.

When the backup is stored on only the IBM Spectrum Protect server, IBM Spectrum Protect Snapshot provides the **Restore Into** option. When the backup is stored on only the local disk, the **Restore Into** option is not available.

After you select **Restore Into**, in the window that is displayed, specify a target location for the restore operation. The target location must be a drive letter or mount point. The size of the target location must be equal to the size of the original volume.

- c) Click **Restore to Point-in-Time** to specify a point in time in the past from which to restore the last version of a volume or mount point.

When this action is selected, you are prompted to provide the following information:

- **PITDate**

Enter the date to establish a point in time to restore a specific version of your custom application or file system backup. Objects that are backed up on or before the date and time that you specify, and that are not deleted before the date and time you specify, are processed. Backup versions that you created after this date and time are ignored.

- **PITTime**

Use this option with the **PITDate** option to establish a point in time to restore a specific version of your custom application or file system backup. Objects that are backed up on or before the date and time that you specify, and objects that are not deleted before the date and time that you specify, are processed. Backup versions that you created after this date and time are ignored. This option is ignored if you do not specify **PITDate** option.

7. To view the status of the restore operation, click **Task List** in the lower half of the results pane. Click **Task Details** to view detailed status information.

Implementing custom application and file system restore scenarios

Apply different restore strategies depending on your network traffic requirements and restore schedule. Follow these typical restore scenarios: restore the entire volume from local shadow volumes, restore individual files from a snapshot or from IBM Spectrum Protect server, or restore the entire volume from VSS image sent to IBM Spectrum Protect server pools.

Restoring the entire volume from a custom application and file system data VSS backup on local shadow volumes

You can use a VSS instant restore operation to restore a local VSS backup and overwrite the entire volume data. You can use a VSS fast restore operation to restore a local VSS backup that copies the files on the volume at the time of the backup.

Procedure

1. If the custom application is running, stop or suspend it. For a file system, close any open handles to that file system.
2. Restore from a VSS backup by specifying the **fcmlcli restore** command with the **/INSTANTRestore=yes** parameter (VSS instant restore) or **/instantrestore=no** parameter (VSS fast restore).

When you restore a VSS backup from an IBM Spectrum Protect server, if the **INSTANTRESTORE** parameter is set to FALSE or **FASTRESTORE** parameter is set to NO, the setting is ignored. The restore operation completes with an image-level restore.

Examples:

Task

By using a VSS instant restore operation, restore the local VSS backup and overwrite the entire volume data without a prompt.

Command: `fcmlcli.exe restore G: /INSTANTRestore=Yes /NOPROMPT`

Task

By using a VSS fast restore operation, restore the local VSS backup that copies the files on the volume at the time of the backup.

Command: `fcmlcli.exe restore G: /INSTANTRestore=No`

3. Complete any actions that are required to achieve a correct state of the custom application and file system files.
4. Restart or resume the custom application.

Restoring individual files from a custom application and file system data VSS backup on local shadow volumes

You can restore files from a VSS backup of custom applications and file system data. The backup is stored on a local shadow volume.

Procedure

1. Mount the VSS backup.

Task

Mount the local VSS backup from drive letter G: to drive letter M:

Command: `fcmdi.exe mount backup G:=M:`

2. If the custom application is running, stop or suspend it.
3. Issue the Windows COPY or XCOPY command or use a tool, for example, Windows Explorer, to copy the files from the VSS backup to your preferred location.

Task

Issue the Windows COPY command to copy the `config.txt` file from drive letter M: to drive letter G: :

Command: `copy M:\config.txt G:\config.txt /y`

4. Complete any actions that are required to achieve a correct state of the custom application.
5. Restart or resume the custom application.
6. Unmount the VSS backup.

Restoring the image backup of a custom application or file system data from IBM Spectrum Protect server

You can restore files from a VSS backup of custom applications and file system data. The backup is stored on an IBM Spectrum Protect server.

Procedure

1. If the custom application is running, stop or suspend it.
2. Restore from an IBM Spectrum Protect backup that was created on an IBM Spectrum Protect server by specifying the `fcmdi restore` command with the `/backupdestination=TSM`.

Task

Restore the backup from drive letter G: volume that was created on an IBM Spectrum Protect server

Command: `fcmdi.exe restore G: /backupdestination=TSM`

3. Restart or resume the custom application.

Restoring an entire volume from a VSS image sent to IBM Spectrum Protect server pools

You can restore a volume from a VSS image that is stored on IBM Spectrum Protect server pools. During backup operations, data is sent as an image backup of the VSS snapshot. The restore operation is a volume-level restore.

Before you begin

- You cannot restore an image-level backup to the volume where the IBM Spectrum Protect backup-archive client is running. To avoid an error, install the IBM Spectrum Protect backup-archive client on the system drive. The same type of failure can occur if you create an application database, for example, a SQL database, under the volume that is being restored.

- For Microsoft VSS operations to succeed, ensure that the file system is of type NTFS or ReFS. You cannot use file systems of type FAT, FAT32, and RAW.

About this task

You can complete this procedure when the **/BACKUPDESTINATION** parameter is set to either TSM or BOTH options.

Procedure

1. If the custom application is running, stop or suspend it.
2. If you plan to restore data from a file system, close open handles to the file system.
3. To restore from a VSS backup that is sent to IBM Spectrum Protect, enter the **fcmlcli restore** command with the **/BACKUPDESTINATION=TSM** parameter.

Archiving the backup to tape with third-party software

You can mount, unmount, and query backup from remote systems. In addition, you can give a user the ability to view, recover, and allow third-party software to access files from the backup. To complete these tasks, the IBM Spectrum Protect Snapshot command-line interface is used to mount the VSS snapshots to servers. When you allow third-party software to archive the backup to tape, the following scenario can be used:

1. **Server A** has a Microsoft Exchange Server database and IBM Spectrum Protect Snapshot installed. IBM Spectrum Protect Snapshot is running on **Server A** to complete regular backups.
2. **Server B** has IBM Spectrum Protect Snapshot installed, along with a third-party tape archive utility. Using the mount command with remote options, **Server B** can be directed to remotely mount the backups that are created by **Server A**. The third-party tape archive utility archives the backups to tape.

Related tasks

[“Troubleshooting file system and custom application VSS restores from IBM Spectrum Protect server” on page 174](#)

File system and custom application VSS restores from IBM Spectrum Protect server are volume image-level restore operations. This type of restore operation might cause the shadow copies, which are created with a system provider for the volume that is being restored, to become invalid and be deleted.

Mounting VSS snapshots to remote servers

You can mount VSS snapshots to remote servers that other users can access. When you remotely mount a snapshot, you must use hardware that includes a feature for creating transportable snapshots. If you use a mix of hardware, all hardware must include a feature for creating transportable snapshots.

Before you begin

You must enable the **Import VSS snapshots only when needed** configuration option before you create a snapshot. Enable this option by clicking **Properties > VSS options**.

To use the remote management and remote mount features, you must install and enable Windows PowerShell 3.0. To download, install and enable the software, follow the instructions in: Microsoft Windows Management Framework 3.0 Downloads (<http://www.microsoft.com/en-s/download/details.aspx?id=34595>)

About this task

When you select the Import VSS snapshots only when needed option, if the VSS hardware provider does not support transportable snapshots, or, if no hardware provider is available, the snapshot completes, but the VSS snapshot is imported and is not transportable.

Restriction: Do not create snapshots that contain a mixture of hardware providers because the snapshot operations fail.

The following steps are based on two servers, Server A and Server B, both with IBM Spectrum Protect Snapshot installed, and attached to the same storage device. The snapshots were created by using IBM Spectrum Protect Snapshot on Server A.

The Command-Line Interface (CLI) and cmdlet methods in the following procedure detail the commands for Data Protection for Microsoft SQL Server.

Procedure

1. Query the snapshots that are available on Server A from Server B, by using one of the following methods:

- GUI method:

Configure Server A as a remote computer and browse the snapshots that are available by using the **Microsoft Management Console (MMC) > Remote Management** feature.

- CLI method:

Configure Server A as a remote computer by using the CLI with the remote computer options.

```
tdpsqlc query fcm * /remotecomputer=ServerA
```

```
/remotecomputeruser=MyDomain\MyUser
```

```
/remotecomputerpassword=MyPassword
```

where *MyDomain* specifies the domain of the remote computer and *MyPassword* specifies the password to the remote computer.

- Cmdlet method:

Configure the remote computer by using the different methods in Windows PowerShell:

```
"Enter-PSSession -ComputerName <Remote Computer>"
```

or

```
"New-PSSession -ComputerName <Remote Computer> -Credential domain\User"
```

or

```
"invoke-command -ComputerName <Remote Computer>
```

```
- credential $(Get-Credential) - ScriptBlock{script}"
```

where *ComputerName* specifies the remote computer.

Tip: From within the Windows Power[®] Shell session, or within the script block script, you can import the IBM Spectrum Protect Snapshot cmdlet modules (ipmo <FlashCopyManager installed path> \FmModuleSql.dll), and then run the "Get-DpSQLBackup" cmdlet.

```
{ipmo 'C:\Program Files\Tivoli\FlashcopyManager\FmModuleSQL.dll';
```

```
Get-DpSqlBackup}
```

2. Mount a snapshot to a remote server, by using one of the following methods:

- CLI method:

From Server B, run the CLI **mount** command with the remote computer options:

```
tdpsqlc mount backup "mydb1=d:\dir1\mntPt" /remotecomputer=ServerA  
/remotecomputeruser=MyDomain\MyUser /remotecomputerpassword=MyPassword
```

where *MyDomain* specifies the domain of the remote computer and *MyPassword* specifies the password to the remote computer.

- Cmdlet method:

From Server B, run the **Mount-DPSQLBackup** cmdlet with the remote computer options:

```
Mount-DPSQLBackup "mydb1=d:\dir1\mntPt" -remotecomputer ServerA  
  
-remotecomputeruser MyDomain\MyUser  
  
-remotecomputerpassword MyPassword
```

where *MyDomain* specifies the domain of the remote computer and *MyPassword* specifies the password to the remote computer.

The active snapshot of the "mydb1" database is mounted from Server A to the d:\dir1\mntPt directory on Server B.

3. Unmount a snapshot, by using one of the following methods:

- CLI method:

From Server B, run the CLI **unmount** command with the remote computer options:

```
tdpsqlc unmount backup "d:\dir1\mntPt" /remotecomputer=ServerA /  
remotecomputeruser=MyDomain\MyUser /remotecomputerpassword=MyPassword
```

where *MyDomain* specifies the domain of the remote computer and *MyPassword* specifies the password to the remote computer.

- Cmdlet method:

From Server B, run the **Dismount-DPSQLBackup** cmdlet with the remote computer options:

```
Dismount-DpSQLBackup "d:\dir1\mntPt" -remotecomputer ServerA -  
remotecomputeruser MyDomain\MyUser -remotecomputerpassword MyPassword
```

where *MyDomain* specifies the domain of the remote computer and *MyPassword* specifies the password to the remote computer.

What to do next

When you use the CLI method to mount VSS snapshots to remote servers, the following restrictions apply:

- On the server where the snapshot is created, no fast restore operation of the full snapshot is available
- Snapshots are mounted with read-only access
- After a snapshot is mounted, the snapshot cannot be mounted to a different location at the same time

When you mount a snapshot remotely and then you delete the snapshot, the state of the mount point varies. The state of the mount point depends on the VSS hardware provider and storage device that you used. When you mount a snapshot remotely, you can delete the snapshot.

When you create a local persistent VSS snapshot, a source and target volume relationship is created. The local persistent VSS snapshot is created on the storage device. In this scenario, when a remote mount operation occurs, the target volume is imported and mounted to the server that sends the request for the remote operation. During the deletion of a snapshot, the snapshots and the relationship between the source and target volumes on the storage device are also deleted. However, the target volume that is

imported and mounted might continue to exist. In addition, the target volume might not be available to the server where the remote mount occurred.

When you use the remote mount feature, consider the following common snapshot deletion scenarios:

- When you manually delete a remotely mounted snapshot. You can use the IBM Spectrum Protect Snapshot software to delete any snapshot by using the CLI or the Microsoft Management Console.
- When a remotely mounted snapshot exceeds snapshot conditions. When either the maximum number of snapshot versions or the maximum number of days to retain a snapshot (as specified by the IBM Spectrum Protect Snapshot policy) is exceeded, the associate snapshot is expired and deleted.

When you enter a mount or query command with the /remotecomputer option, enable CLI tracing to debug the problem, if a problem occurs. For tracing, append /tracefile=filename.trc /traceflag=service to the command.

After tracing is enabled, the CLI generates trace files for the local and remote systems. On the local system, you can view the file that you specified. In addition, on the local and remote systems, a trace file is also created. This file has the same name as the file stored on the local system and the file name concludes with the following suffix appended to the file type extension: `_remote`.

In addition to the CLI trace file, you can enable tracing on the agent, enable tracing on both the local and remote systems, and collect trace files on both the local and remote systems.

Related tasks

[Gathering trace and log files for remote systems](#)

Collecting diagnostic data for a remote system, by using IBM Spectrum Protect Snapshot, is different to collecting data for a local system. You can update the **Diagnostics** property page to collect the correct log and trace files for remote systems.

[Automating Microsoft Exchange Server tasks](#)

You can automate a workload by entering Windows PowerShell cmdlets or commands in the integrated command-line interface.

Enabling Windows PowerShell Remoting for Remote Management and Remote Mounting

From a single IBM Spectrum Protect Snapshot installation, you can manage all of the IBM Spectrum Protect Snapshot installations in your organization.

Before you begin

To use the remote management features, you must have the minimum required Windows Powershell and Windows management framework installed. For more information, see the IBM Spectrum Protect Snapshot V8.1.7 hardware and software requirements: <https://www-01.ibm.com/support/docview.wss?uid=ibm10792377>.

Procedure

Enabling Windows PowerShell Remoting is a task outside the scope of this documentation. For reference, the following PowerShell cmdlets are provided.

1. Enable remote management for IBM Spectrum Protect Snapshot installations or the Remote Mounting feature by entering the following Windows PowerShell cmdlets.

```
Enable-PSRemoting -force
```

- a) Add the IBM Spectrum Protect Snapshot servers to the trusted hosts list by entering the following command on each remote system:

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value  
remote_server_name -Force -Concatenate
```

where *remote_server_name* specifies the remote server.

- b) Restart the winrm service by entering the following command:

```
Restart-Service winrm
```

2. Optional: If you use the IBM Spectrum Protect Snapshot for Microsoft Exchange Server software, enable the Windows PowerShell Remoting feature with Credential Security Support Provider (CredSSP) authentication. Complete the following steps:

- a) On the primary system, enter the following command to enable the Windows PowerShell Remoting feature with CredSSP:

```
enable-wsmancredssp -role client -delegatecomputer remote_computer_name
```

where *remote_computer_name* specifies the remote computer.

- b) On each remote system that runs the IBM Spectrum Protect Snapshot for Microsoft Exchange Server software, enter the following command to enable the Windows PowerShell Remoting feature with CredSSP:

```
enable-wsmancredssp -role server
```

3. Verify that the Windows PowerShell Remoting feature is configured by using one of the following methods:

- Use the "Test-WSMan" cmdlet to test whether the WinRM service is running on the remote computer.
 - a. On the primary system, enter the following cmdlet to verify that the Windows PowerShell Remoting feature is configured correctly:

```
Test-WSMan remote_server_name
```

where *remote_server_name* specifies the remote server.

- b. On the remote system, enter the following cmdlet to verify that the Windows PowerShell Remoting feature is configured correctly:

```
Test-WSMan primary_server_name
```

where *primary_server_name* specifies the primary server.

Restriction: For the remote mount feature, you must use the same computer name that you use with the /RemoteComputer CLI option.

- To verify that the Windows PowerShell Remoting feature is configured, enter the following cmdlets:
 - a. On the primary system and the remote system, enter the following cmdlet:

```
invoke-command -computername remote_server_name
```

```
-scriptblock {pwd} -Credential $creds
```

where *remote_server_name* specifies the remote server.

- b. On the primary system and the remote system when Credential Security Support Provider (CredSSP) authentication is enabled, enter the following cmdlet:

```
invoke-command -computername remote_server_name
```

```
-scriptblock {pwd} -Credential $creds -Authentication Credssp
```

where *remote_server_name* specifies the remote server.

Restriction: For the remote mount feature, you must use the same computer name that you use with the /RemoteComputer CLI option. In addition, when you use the CLI command for the remote mount feature, use the same user name and password that you use with the /RemoteComputerUser and /RemoteComputerPassword CLI options.

Adding remote systems

To remotely manage installations, complete the following steps to add remote systems.

Procedure

1. From **Microsoft Management Console (MMC) > IBM Spectrum Protect Snapshot**, in the **Actions** pane, click **Manage Computers**.

2. Verify that the local system is listed in both the **Tree Nodes** and **Computers** panes.

3. From the **Tree Nodes** pane, click the add icon.

The icon is green and resembles the symbol for addition.

4. Type the name and description for the new tree node.

5. From the **Computers** pane, click the add icon.

The computers that you add are associated with the tree node that you are creating. If you add only one computer, the tree node type can be either **Dashboard** or **Group**. If you add more than one computer, the tree node type is **Group**. If you add only one computer, from the **Tree Nodes** pane, you can toggle between the **Dashboard** and **Group** types.

6. Type the system name and a description. For systems that are not in the domain, provide the fully qualified address.

Alternatively, to select a system that is based on another system in the domain or to read a list of computers from a file, on the **Computers** pane, click **Import**. Clicking **Import** displays a dialog called **Add Computers**. From the **Add Computers** dialog, there are two tabs: Active Directory and Import. To complete the **Add Computers** dialog window entries, complete the following steps:

a) For the **Active Directory** tab, complete these fields

Domain

The current domain is displayed. The domain cannot be changed.

Location

The organizational unit that is used to search for computers. The default value is displayed.

Name

By default, the wildcard character (*) is displayed. You can leave the default value or enter a specific name.

Account

The current account is displayed. If you want to use a different account, click **Search** to search the domain for other computers. The Search capability is enabled only when the **Location** and **Name** fields have values.

b) For the **Import** tab, browse to find a comma-separated values (.CSV) file that contains computer entries. After you find a .CSV file and click **Import**, the contents of the .CSV file are read as entries and are added to the list.

The following .CSV file is an example of a valid .CSV file for the import activity:

```
NewNode1,Group1,CurrentUser,Test node 1
NewNode2
NewNode3,,Description of NewNode3
NewNode4,Group2,CurrentUser,Test node 4
```

The first column (the node name) is required. The other data columns are optional. The list is processed by position. For the group, if a group does not exist, the group is created.

7. From the **Computers** pane, click **Test Connection**.

The test status is reported in the Message column of the **Computers** table.

8. Click **OK** to close the **Manage Computers** window.

9. Verify that the tree node is listed in the navigation tree.

The remote node does not have all of the functionality available for local systems. For example, entries for learning, online support, and favorite links are not displayed.

For tree node type **Dashboard**, the main window displays the **Protect**, **Recover**, and **Automate** tabs. For tree node type **Group**, the main window displays the **Group Dashboard**, **Group Reports**, and **Group Commands** tabs.

10. After you add systems, you can remove (delete) the systems. You can also select the system to edit the properties, including tree node type, that you entered when you added the system. If you want to change the order of the systems that are displayed in the navigation tree, from the **Manage Computers** window there are GUI controls that you can use to change the order.

Viewing, printing, and saving reports

You can access reports on recent activity and historical managed capacity. You can determine which licenses and software are installed.

Procedure

1. Select **Reporting** in the **Manage** section.
A list of available reports is displayed. Each report provides a summary of the report contents.
2. Select a report from the list.
The selected report displays.
3. To print or save the current report, click the appropriate icon at the top of the report.

Generating group reports

When you use the **Group** tree nodes in **Management Console** (MMC) to create a group, the **Group Dashboard**, **Group Reports**, and **Group Commands** tabs replace the **Protect**, **Recover**, and **Automate** tabs.

Before you begin

Your system must run Windows 2008 or later versions, PowerShell 3.0 or later, and IBM Spectrum Protect Snapshot. Workloads and backed up data must be configured successfully.

Procedure

1. In MMC, select the group with the added systems.
2. In the main window, select the **Group Reports** tab.
The list of reports is displayed.
3. In the **Actions** pane, verify that the group name is correct.
4. In the **Reports** section, click **Refresh** to refresh the data that is displayed.

Chapter 6. Automating tasks

With IBM Spectrum Protect Snapshot *automation* capability, you can run commands from the command line, create scripts, schedule tasks, and use Microsoft Management Console (MMC) to start tasks. The tasks that you can automate are based on the scripts and schedules that you create.

IBM Spectrum Protect Snapshot supports you automating tasks from the command-line interface or Microsoft Windows PowerShell command prompt (Version 3.0 or later). You can also use the **Automate** tab in the MMC.

Preparing to use Windows PowerShell cmdlets

IBM Spectrum Protect Snapshot includes a set of Windows PowerShell cmdlets to help you manage Exchange Server, SQL Server, file system, and custom application data in your environment.

About this task

You can use the cmdlets that are provided with IBM Spectrum Protect Snapshot in Windows environments.

IBM Spectrum Protect Snapshot cmdlets support a seamless management environment and greatly improve remote management and automation capabilities. You can aggregate cmdlets together to form commands and use the large volume of existing cmdlets from other vendors.

Before you use the cmdlets, complete the following steps.

Procedure

1. Log on to the system as an administrator.
2. From a Windows PowerShell command line, enter the following command:

```
set-executionpolicy remotesigned
```

3. During installation of IBM Spectrum Protect Snapshot, the following Windows PowerShell modules are imported automatically from the FlashCopyManager folder.

- FmModuleExc.dll
- FmModuleFs.dll
- FmModuleMMC.dll
- FmModuleSQL.dll

If you wish to import the Windows PowerShell modules manually, from the Windows PowerShell command prompt, import modules, with the administrator credentials, as follows:

- a) Go to the FlashCopyManager folder.
- b) Enter the following commands:

```
import-module .\FmModuleExc.dll
import-module .\FmModuleFs.dll
import-module .\FmModuleMMC.dll
import-module .\FmModuleSQL.dll
```

- c) (Optional) To use the cmdlets in these modules any time that you start Windows PowerShell, add the following lines to your profile.

The following path is the default profile path.

```
$path = (get-itemproperty -path "HKLM:\SOFTWARE\IBM\FmModuleManager\
currentversion\mmc" -ea SilentlyContinue).path
if ($null -ne $path)
{
```

```
dir "$path\fmmodule*.dll" | select -expand fullname | import-module
    -force -Global
}
```

What to do next

For information about creating, running, monitoring, and troubleshooting scripts with cmdlets, see Windows PowerShell 3.0 or later documentation. For more information about Windows PowerShell cmdlets, consistent naming patterns, parameters, arguments, and syntax, see this web page as a starting point: [Microsoft TechNet: Getting Started with Windows PowerShell \(http://technet.microsoft.com/en-us/library/hh857337.aspx\)](http://technet.microsoft.com/en-us/library/hh857337.aspx).

Cmdlets for Microsoft Management Console

The following list identifies the cmdlets that you can use when interacting with Microsoft Management Console (MMC).

- **Clear-FcmMmcManagedCapacityHistory**
- **Clear-FcmMmcScheduledActivityHistory**
- **Disable-FcmMmcSchedule**
- **Enable-FcmMmcSchedule**
- **Get-FcmMmcActivity**
- **Get-FcmMmcComputerInformation**
- **Get-FcmMmcManagedCapacityHistory**
- **Get-FcmMmcReport**
- **Get-FcmMmcSchedule**
- **Get-FcmMmcScheduledActivity**
- **New-FcmMmcSchedule**
- **Remove-FcmMmcSchedule**
- **Set-FcmMmcSchedule**
- **Start-FcmMmcSchedule**

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help New-FcmMmcSchedule
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help New-FcmMmcSchedule -examples
```

For more information, enter:

```
get-help New-FcmMmcSchedule -detailed
```

For technical information, enter:

```
get-help New-FcmMmcSchedule -full
```

For online product information, enter:

```
get-help New-FcmMmcSchedule -online
```

For information about a specific parameter, enter:

```
help New-FcmMmcSchedule -Parameter backupdestination
```

To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Cmdlets for protecting Microsoft Exchange Server data

The following table identifies the cmdlets that you can use to protect Microsoft Exchange Server data.

Table 34. Cmdlets to protect Microsoft Exchange Server data. The following table identifies the cmdlets that you can use to protect Microsoft Exchange Server data.

Cmdlet name	Related command-line interface command	Short description
Add-DpExcPolicy	tdpexcc create policy	Create a policy for IBM Spectrum Protect Snapshot for Exchange Server.
Backup-DpExcComponent	tdpexcc backup	Back up a Microsoft Exchange database.
Copy-DpExcPolicy	tdpexcc copy policy	Copy an existing policy.
Dismount-DpExcBackup	tdpexcc unmount backup	Dismount a backup.
Get-DpExcBackup	tdpexcc query tsm *	Query backups.
Get-DpExcComponent	tdpexcc query exchange	Query the Exchange Server for all databases that are available for backup.
Get-DpExcConfig	tdpexcc query tdp	Display configuration information.
Get-DpExcConnection	tdpexcc query tsm	Query a list of the current values set in the configuration file for IBM Spectrum Protect.
Get-DpExcInformation	tdpexcc query exchange	Query general local Exchange Server information.
Get-DpExcMailboxLocationHistory	tdpexcc q tsm / showMailboxInfo	Query the mailbox location history.
Get-DpExcManagedCapacity	tdpexcc query managedcapacity	Query managed capacity for Microsoft Exchange Server.
Get-DpExcPolicy	tdpexcc query policy	Display policy information.
Mount-DpExcBackup	tdpexcc mount backup	Mount a backup to provide access to the files that the backup contains. You can mount a read-only or a read-write backup.
Remove-DpExcBackup	tdpexcc delete backup	Remove the backup.
Remove-DpExcPolicy	tdpexcc delete policy	Delete the policy.
Reset-DpExcTsmPassword	tdpexcc changetspmpassword	Change the IBM Spectrum Protect password that is used by IBM Spectrum Protect Snapshot for Exchange Server.
Restore-DpExcBackup	tdpexcc restore	Restore a backup.
Restore-DpExcMailbox	tdpexcc restore mailbox	Restore a mailbox. You can mount a read-only or a read-write backup to perform the mailbox restore.
Set-DpExcConfig	tdpexcc set paramname	Set the application configuration parameters in a configuration file.

Table 34. Cmdlets to protect Microsoft Exchange Server data. The following table identifies the cmdlets that you can use to protect Microsoft Exchange Server data. (continued)

Cmdlet name	Related command-line interface command	Short description
Set-DpExcPolicy	tdpexcc update policy	Update a policy.

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help Backup-DpExcComponent
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help Backup-DpExcComponent -examples
```

For more information, enter:

```
get-help Backup-DpExcComponent -detailed
```

For technical information, enter:

```
get-help Backup-DpExcComponent -full
```

For online product information, enter:

```
get-help Backup-DpExcComponent -online
```

For information about a specific parameter, enter:

```
help Backup-DpExcComponent -Parameter backupdestination
```

To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Cmdlets for protecting Microsoft SQL Server data

The following table identifies the cmdlets that you can use to protect Microsoft SQL Server data.

Table 35. Cmdlets to protect Microsoft SQL Server data. The following table identifies the cmdlets that you can use to protect Microsoft SQL Server data.

Cmdlet name	Related command-line interface command	Short description
Add-DpSqlPolicy	tdpsqlc create policy	Create a policy for Microsoft SQL Server data.
Backup-DpSqlComponent	tdpsqlc backup	Backup SQL components.
Copy-DpSqlPolicy	tdpsqlc copy policy	Copy an existing policy to a new policy.
Dismount-DpSqlBackup	tdpsqlc unmount backup	Dismount a backup.
Dismount-DpSqlClone	not applicable	Unmounts the clone from the server.
Get-DpSqlBackup	tdpsqlc query tsm *	Query the backups that are stored on the server.
Get-DpSqlClone	not applicable	Queries the clones that are available.
Get-DpSqlComponent	tdpsqlc query sql *	Query the databases that are available on the SQL Server.
Get-DpSqlConfig	tdpsqlc query tdp	Display configuration information.

Table 35. Cmdlets to protect Microsoft SQL Server data. The following table identifies the cmdlets that you can use to protect Microsoft SQL Server data. (continued)

Cmdlet name	Related command-line interface command	Short description
Get-DpSqlConnection	tdpsqlc query tsm	Display the IBM Spectrum Protect API and server information.
Get-DpSqlFileGroups	not applicable	Display all file and group information about specified SQL Server databases.
Get-DpSqlInformation	tdpsqlc query sql	Display specified SQL Server information.
Get-DpSqlManagedCapacity	tdpsqlc query managedcapacity	Assist with storage planning by determining the amount of managed capacity that is in use.
Get-DpSqlPolicy	tdpsqlc query policy	Query policy.
Mount-DpSqlBackup	tdpsqlc mount backup	Mount a backup that provides access to the files that are contained by the backup. You can mount a backup as read-only or read/write.
Mount-DpSqlClone	not applicable	Mounts the created clone to any server. Clones are always mounted as read/write.
New-DpSqlCloneFromComponent	not applicable	Creates a persistent VSS snapshot from a production database to be used as a clone.
Remove-DpSqlBackup	tdpsqlc delete backup and tdpsqlc deactivate	Delete a VSS backup of an SQL Server database, or deactivate one or more active legacy backup objects on the IBM Spectrum Protect server.
Remove-DpSqlClone	not applicable	Deletes the clone from the SQL server.
Remove-DpSqlPolicy	tdpsqlc delete policy	Delete a local policy.
Reset-DpSqlTsmPassword	tdpsqlc changetsmpassword	Change the IBM Spectrum Protect password that is used by IBM Spectrum Protect Snapshot for SQL Server.
Restore-DpSqlBackup	tdpsqlc restore	Restore backups of Microsoft SQL Server data.
Set-DpSqlConfig	tdpsqlc set paramname	Set the IBM Spectrum Protect Snapshot for SQL Server configuration parameters in the configuration file.
Set-DpSqlPolicy	tdpsqlc update policy	Change an existing policy.

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help Get-DpSqlBackup
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help Get-DpSqlBackup -examples
```

For more information, enter:

```
get-help Get-DpSqlBackup -detailed
```

For technical information, enter:

```
get-help Get-DpSqlBackup -full
```

For online product information, enter:

```
get-help Get-DpSqlBackup -online
```

For information about a specific parameter, enter:

```
help Get-DpSqlBackup -Parameter backupdestination
```

To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Cmdlets for protecting custom application and file system data

The following table identifies the cmdlets that you can use to protect custom application and file system data.

Cmdlet name	Related command-line interface command (if available)	Short description
Add-FcmFsPolicy	fccli create policy	Add a VSS policy binding statement.
Add-FcmFsVssPolicy	fccli insert vsspolicy	Insert a new VSS policy binding statement.
Backup-FcmFsComponent	fccli backup	Create a VSS snapshot backup of volumes and mount points.
Copy-FcmFsPolicy	fccli copy policy	Copy a policy.
Dismount-FcmFsBackup	fccli unmount backup	Unmount a mounted backup.
Get-FcmFsBackup	fccli query backup	Display information about the backup.
Get-FcmFsComponent	fccli query component	Query the VSS components that are available on the system.
Get-FcmFsConfig	fccli query config	Display configuration information.
Get-FcmFsConnection	fccli query config	Query IBM Spectrum Protect server connection information.
Get-FcmFsManagedCapacity	fccli query managedcapacity	Assist with storage planning by determining the amount of managed capacity that is in use.
Get-FcmFsPolicy	fccli query policy	Display policy information.
Get-FcmFsVSSPolicy	fccli query vsspolicy	Return the VSS policy binding statements that are stored in the configuration file.

Table 36. Cmdlets to protect custom application and file system data. The following table identifies the cmdlets that you can use to protect custom application and file system data. (continued)

Cmdlet name	Related command-line interface command (if available)	Short description
Mount-FcmFsBackup	fcmdi mount backup	Mount a backup that provides access to the files that the backup contains. You can mount a read-only or a read-write backup.
Remove-FcmFsBackup	fcmdi delete backup	Delete a backup from IBM Spectrum Protect Snapshot storage.
Remove-FcmFsPolicy	fcmdi delete policy	Remove a policy.
Remove-FcmFsVssPolicy	fcmdi delete vsspolicy	Delete a VSS policy binding statement.
Reset-FcmFsTsmPassword	fcmdi changetsmpassword	Change the IBM Spectrum Protect password that is used by the IBM Spectrum Protect Snapshot for File Systems and Custom Applications.
Restore-FcmFsBackup	fcmdi restore	Restore a backup.
Set-FcmFsConfig	fcmdi update config	Update configuration for file systems and custom applications.
Set-FcmFsPolicy	fcmdi update policy	Update an existing policy.
Set-FcmFsVssPolicy	fcmdi update vsspolicy	Update an existing VSS policy binding statement.

To view the details about a specific cmdlet, run the **Get-Help** cmdlet with the cmdlet name. For example:

```
Get-Help Backup-FcmFsComponent
```

To continue the example, to see examples for the cmdlet, enter:

```
get-help Backup-FcmFsComponent -examples
```

For more information, enter:

```
get-help Backup-FcmFsComponent -detailed
```

For technical information, enter:

```
get-help Backup-FcmFsComponent -full
```

For online product information, enter:

```
get-help Backup-FcmFsComponent -online
```

For information about a specific parameter, enter:

```
help Backup-FcmFsComponent -Parameter backupdestination
```

To display the help in a separate window, include the **-showwindow** parameter with the **help** command.

Automating Microsoft Exchange Server tasks

You can automate a workload by entering Windows PowerShell cmdlets or commands in the integrated command-line interface.

About this task

An integrated command line is available in the task window from which you can enter PowerShell cmdlets or command-line interface commands. You use the **Automate** view to work with commands. You can create, save, store, and schedule commands to run at the scheduled time.

Procedure

1. To open the **Automate** view, select a workload that you want to work with and click **Automate**.
2. Change **PowerShell** to **Command Line**.
3. To run a command, type a command in the details pane and click the **Execute** icon.

You can enter the commands with or without specifying `fcmdi`.

For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

```
q fcm
```

You can also run a saved task by clicking the **Open** icon, selecting the command file, and clicking the **Execute** icon.

The output is displayed in the main window.

4. Click the **Save** icon and follow the prompts to save a command for future use.
5. To schedule a command, click the **Schedule this command** icon to open the scheduling wizard. Follow the prompts in the wizard to create a schedule for the command.
The output of the command is displayed in the results pane.
6. (Optional) Save or send the command output to an email address.

What to do next

You can automate commands from the **Protect**, **Recover**, **Schedule**, and **Task List** views in Microsoft Management Console (MMC):

1. Start MMC and select a workload in the navigation tree.
2. Click the tab for the task you want to do (**Protect** or **Recover**).
3. Automate the command by using one of the following methods:

Result pane

Select the item for your task in the result pane, and select **Run Scheduled** in the toolbar menu. In the **Actions** pane, click the appropriate task. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

Task List pane

When a task is submitted, it displays in the task list pane. Select the appropriate task, then click **Schedule command script** in the task list toolbar. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

You can also right-click a task in the Task List pane and click **Copy**. Then, click the **Automate** tab and paste the command in the field.

Automating Microsoft SQL Server tasks

You can automate a workload by entering Windows PowerShell cmdlets or commands in the integrated command-line interface.

About this task

An integrated command line is available in the task window from which you can enter PowerShell cmdlets or command-line interface commands. You use the **Automate** view to work with commands. You can create, save, store, and schedule commands to run at the scheduled time.

Procedure

1. To open the **Automate** view, select a workload that you want to work with and click **Automate**.
2. Change **PowerShell** to **Command Line**.
3. To run a command, type a command in the details pane and click the **Execute** icon.

You can enter the commands with or without specifying `fcmdi`.

For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

```
q fcm
```

You can also run a saved task by clicking the **Open** icon, selecting the command file, and clicking the **Execute** icon.

The output is displayed in the main window.

4. Click the **Save** icon and follow the prompts to save a command for future use.
5. To schedule a command, click the **Schedule this command** icon to open the scheduling wizard. Follow the prompts in the wizard to create a schedule for the command.
The output of the command is displayed in the results pane.
6. (Optional) Save or send the command output to an email address.

What to do next

You can automate commands from the **Protect**, **Recover**, **Schedule**, and **Task List** views in Microsoft Management Console (MMC):

1. Start MMC and select a workload in the navigation tree.
2. Click the tab for the task you want to do (**Protect** or **Recover**).
3. Automate the command by using one of the following methods:

Result pane

Select the item for your task in the result pane, and select **Run Scheduled** in the toolbar menu. In the **Actions** pane, click the appropriate task. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

Task List pane

When a task is submitted, it displays in the task list pane. Select the appropriate task, then click **Schedule command script** in the task list toolbar. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

You can also right-click a task in the Task List pane and click **Copy**. Then, click the **Automate** tab and paste the command in the field.

Automating custom applications and file system tasks

You can automate a workload by entering Windows PowerShell cmdlets or commands in the integrated command-line interface.

About this task

An integrated command line is available in the task window from which you can enter PowerShell cmdlets or command-line interface commands. You use the **Automate** view to work with commands. You can create, save, store, and schedule commands to run at the scheduled time.

Procedure

1. To open the **Automate** view, select a workload that you want to work with and click **Automate**.
2. Change **PowerShell** to **Command Line**.
3. To run a command, type a command in the details pane and click the **Execute** icon.

You can enter the commands with or without specifying `fcmdi`.

For example, for each selected workload instance, you can enter a single command or multiple commands, such as:

```
q component  
q backup
```

You can also run a saved task by clicking the **Open** icon, selecting the command file, and clicking the **Execute** icon.

The output is displayed in the main window.

4. Click the **Save** icon and follow the prompts to save a command for future use.
5. To schedule a command, click the **Schedule this command** icon to open the scheduling wizard. Follow the prompts in the wizard to create a schedule for the command.
The output of the command is displayed in the results pane.
6. (Optional) Save or send the command output to an email address.

What to do next

You can automate commands from the **Protect**, **Recover**, **Schedule**, and **Task List** views in Microsoft Management Console (MMC) :

1. Start MMC and select a workload in the tree view.
2. Click the tab for the task you want to do (**Protect** or **Recover**).
3. Automate the command by using one of the following methods:

Result pane

Select the item for your task in the result pane, and select **Run Scheduled** in the toolbar menu. Click the appropriate task in the **Action** pane. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

Task List pane

When a task is submitted, it displays in the task list pane. Select the appropriate task, then click **Schedule command script** in the task list toolbar. When the schedule wizard starts, enter the information for each prompt to create a scheduled task.

You can also right-click a task in the Task List pane and click **Copy**. Then, click the **Automate** tab and paste the command in the field.

Scheduling tasks

Automate your data protection with IBM Spectrum Protect Snapshot scheduling. IBM Spectrum Protect Snapshot uses the Windows Scheduler to automate backup and restore operations. You can also schedule tasks by using the TSM Scheduler, and by using the PowerShell cmdlets that are available for use when you interact with Microsoft Management Console (MMC).

Before you begin

The scheduling wizards contain templates for PowerShell and command-line scripts. The default is PowerShell. For more information about the PowerShell cmdlets that are available for scheduling tasks, open a Windows PowerShell prompt and change directories to C:\Program Files\Tivoli\FlashCopyManager. Enter the following command:

```
gcm -mod FmModuleMMC *schedule*
```

You might see a list of available scheduling cmdlets like the following sample:

```
CommandType Name ModuleName
-----
Cmdlet Clear-FcmMmcScheduledActivityHistory FmModuleMMC
Cmdlet Disable-FcmMmcSchedule FmModuleMMC
Cmdlet Enable-FcmMmcSchedule FmModuleMMC
Cmdlet Get-FcmMmcSchedule FmModuleMMC
Cmdlet Get-FcmMmcScheduledActivity FmModuleMMC
Cmdlet New-FcmMmcSchedule FmModuleMMC
Cmdlet Remove-FcmMmcSchedule FmModuleMMC
Cmdlet Set-FcmMmcSchedule FmModuleMMC
Cmdlet Start-FcmMmcSchedule FmModuleMMC
```

About this task

Restriction: With IBM Spectrum Protect Snapshot scheduling operations, you can schedule tasks to run periodically. However, you cannot schedule tasks to run only one time.

Procedure

1. Create and edit new schedules.

Use the Scheduling wizard to guide you through the steps to define a local scheduled data protection task. The Scheduling wizard is available in the **Action** pane.

You can select the type of scheduler you want to use to manage your scheduled operations. Click the relevant check box to select either the local Windows Scheduler or the IBM Spectrum Protect client scheduler. To use the IBM Spectrum Protect client scheduler, you must configure the client to use the client acceptor to manage the scheduler. For more information, see [Configuring the client to use the client acceptor service to manage the scheduler](http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.2/client/c_cfg_dsmcutil_usewin.html)(http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.2/client/c_cfg_dsmcutil_usewin.html).

You can create the following types of schedules:

Hourly

This type of schedule starts at a set time and runs indefinitely or for a set duration. It can be repeated at a specified time. Despite the duration or repeat settings, this type of schedule runs within one 24-hour period only.

Daily

This type of schedule starts at a set time and repeats each day as specified.

Weekly

This type of schedule starts at a set time and repeats every week as specified.

Monthly

This type of schedule starts at a set time and repeats every month as specified.

2. If you select **Windows Scheduler**, then your schedule is managed by the local Windows schedule. If you select **IBM Spectrum Protect client scheduler**, a macro file is created that contains the IBM Spectrum Protect server commands required to create your scheduling task. You can specify IBM Spectrum Protect server details, such as administrator user details and the port and the IBM Spectrum Protect server address. You can also review and edit the created macro file.
3. Select **PowerShell** in the scheduling wizard, and complete one or more of the following tasks:
 - Specify a single schedule to complete workloads as a single scheduled task. For example, you can complete a full backup on Sunday and incremental backups on other days.
 - Select the **MMC template** to generate and email a report. When you select **PowerShell** in the scheduling wizard, four templates are available: file system, SQL Server, Exchange Server, and MMC. The other templates include statements to ensure that the correct working directory is used, and that error information is handled correctly. The templates also include sample statements to run queries and backups.
4. After you define your schedule, run it manually. Select the schedule and, in the **Actions** pane, click **Run**.

For more granular control of your schedules, access the Windows scheduled tasks control pane.

By default, IBM Spectrum Protect Snapshot schedules are activated by using Windows System Account permissions. If a schedule requires different Windows permissions, click **Run as** and enter the appropriate account and password. You cannot specify the percentage (%) character.

All defined schedules are displayed.
5. The scheduled history log file keeps entries for 60 days, by default. To override this default value, change the scheduled history log prune value in the main IBM Spectrum Protect Snapshot settings. In the Tree View, select the computer node that you want, and in the **Actions** pane, click **Properties**.

Chapter 7. Troubleshooting

IBM Spectrum Protect Snapshot supports you in protecting Microsoft Exchange Server and Microsoft SQL Server databases, file systems, and custom applications.

If you encounter a problem, you typically start with a symptom, or set of symptoms, and trace the root cause. Problem determination, however, is not the same as problem solving. During the process of problem determination, you might obtain sufficient information to enable you to solve the problem.

In some cases, you cannot solve a problem even after you determine its cause. For example, a performance problem might be caused by a limitation of your hardware. Debugging a VSS issue might lead you to analyze other components, for example, the VSS hardware provider, the VSS system, or the Multi Path Input Output (MPIO).

Diagnosing problems

One of the most difficult challenges of troubleshooting in a client/server environment is determining which component is the origin of the problem. VSS diagnostic wizards are available to help you test VSS snapshots on your system. You can determine whether the source of the problem is a general VSS issue or a IBM Spectrum Protect issue.

Error log files for IBM Spectrum Protect Snapshot components

If you are protecting an Exchange or SQL Server, or a file system or custom application, an error condition occurs when you are protecting, you can view several log files to help diagnose the problem.

For example, you can confirm that IBM Spectrum Protect Snapshot failed over by searching entries about the secondary server in the following log files:

- For IBM Spectrum Protect Snapshot for Microsoft Exchange Server, check the following files:
 - `Tivoli\tsm\TDPEXchange\dsierror.log`
 - `Tivoli\tsm\baclient\dsmerror.log`
- For IBM Spectrum Protect Snapshot for Microsoft SQL Server, check the following files:
 - `Tivoli\tsm\TDPSQL\dsierror.log`
 - `Tivoli\tsm\baclient\dsmerror.log`
- For IBM Spectrum Protect Snapshot file system and custom applications, check the following files:
 - `Tivoli\flashcopymanager\dsierror.log`
 - `Tivoli\tsm\baclient\dsmerror.log`
- IBM Spectrum Protect Snapshot for Microsoft Exchange Server, and IBM Spectrum Protect Snapshot for Microsoft SQL Server, logs information about backup, restore, and delete commands to the Tivoli Event Console.
- IBM Spectrum Protect Snapshot for Microsoft Exchange Server logs information, by default, to the `tdpexc.log` file in the directory where IBM Spectrum Protect Snapshot for Microsoft Exchange Server is installed. IBM Spectrum Protect Snapshot for Microsoft SQL Server logs information, by default, to the `tdpsql.log` file in the directory where IBM Spectrum Protect Snapshot for Microsoft SQL Server is installed. The log file indicates the date and time of a backup, the data that is backed up, and any error messages or completion codes. This file is important and must be monitored.
- The IBM Spectrum Protect API logs API error information, by default, to the `dsierror.log` file in the directory where IBM Spectrum Protect Snapshot for Microsoft Exchange Server, or IBM Spectrum Protect Snapshot for Microsoft SQL Server, is installed. No backup statistics are contained in this log. The `dsierror.log` file cannot be marked as read-only.

- IBM Spectrum Protect Snapshot for Microsoft Exchange Server logs information to the Exchange Server error log. IBM Spectrum Protect Snapshot for Microsoft SQL Server logs information to the SQL Server error log. The error log information can be viewed using the Exchange Server or SQL Server administration tools.
- The IBM Spectrum Protect scheduler logs information to both the `dsmsched.log` and the `dsierror.log` files. By default, these files are in the directory where the IBM Spectrum Protect backup-archive client is installed.

Note: Output from scheduled commands is sent to the scheduler log file (`dsmsched.log`). After the scheduled work completes, check the log to ensure that the work completed successfully.

When a scheduled command is processed, the scheduler log can contain the following entry:

```
Scheduled event eventname completed successfully
```

This entry is merely an indication that IBM Spectrum Protect successfully issued the scheduled command that is associated with the *eventname*. No attempt is made to determine the success or failure of the command. You can assess the success or failure of the command by evaluating the return code from the scheduled command in the scheduler log. The scheduler log entry for the command's return code is prefaced with the following text:

```
Finished command. Return code is: return_code_number
```

- Windows System and Application Event Log.
- For VSS operations, view the `dsmerror.log` file in the backup-archive client installation directory.

Trace files for IBM Spectrum Protect Snapshot components

When you gather trace files for local or remote systems, the files are automatically copied, compressed, and stored in the `C:\Program Files\Tivoli\flashcopymanager\problemdetermination` folder or other information.

MMC

Options are stored in the MMC user settings file. The following file is generated:

```
TraceFm.trc
TraceUx.trc
```

Data Protection

Tracing options are stored in the MMC user settings file and submitted to the Data Protection component as part of the command. The following file is generated:

```
TraceFileFs.trc
TraceFileSql.trc
TraceFileExc.trc
```

Agent

Tracing options are stored in the VSS Requestor `dsm.opt` file. The following file is generated:

```
TraceFileAgent.trc
```

API

Tracing options are stored in the respective Data Protection `dsm.opt` file. The following file is generated:

```
TraceFileFsAPI.trc
```

Diagnosing VSS issues

You can test VSS persistent, non-persistent, and resync snapshots on your system with the assistance of a VSS diagnostics wizard.

Before you begin



Attention: Do not run these tests if you are already using SAN Volume Controller or Storwize V7000 space-efficient snapshots on your computer. If you do so, existing snapshots might be removed.

Procedure

1. Start Microsoft Management Console (MMC).
2. To open the diagnostics wizard, complete these steps:
 - a) Click **Diagnostics** in the results pane of the welcome page.
 - b) In the **Actions** pane, click **VSS Diagnostics**.A list of volumes are displayed, and the status of each test is displayed when it is completed.
3. To view the results of the persistent and non-persistent snapshot testing, complete these steps:
 - a) Select the volumes or mount points to test and click **Next**.
 - b) Click **Show VSS Information** to view details about the VSS providers, writers, and snapshots that are available on your system.

The results of the persistent and non-persistent snapshot testing displays as Passed or Failed.

4. To view the results of the resync snapshot testing, complete these steps:



CAUTION: VSS ResyncLUNs API instant restore tests will revert the data on the volume to an earlier time. Do not enable these instant restore tests on production volumes as data loss may occur.

- a) To test if the selected volumes support the VSS ResyncLuns API, select a volume and then click **Next**.
- b) Verify that the **Testing resync snapshot** field indicates a successful result.

The results of the resync snapshot testing display as Passed or Warning.

Note: On non-IBM storage devices, resync snapshots are necessary only for instant restore. Resync snapshots have no impact on backup and fast restore on non-IBM storage devices.

5. Review the results of the snapshot testing and click **Next**.

The final results of the persistent and non-persistent snapshot testing display as Success or Unsuccessful.
6. Depending on the results, complete these steps:
 - If the testing status is a success, click **Finish** and exit the wizard.
 - If the testing status is not successful, click **Previous** and review information in the **Rule** dialog.
7. Return to the **Management** window and begin backup operations.

Resolving reproducible problems

When a component fails to operate as designed, try to reproduce the problem and capture information about the current operating environment at the time of the error. You can troubleshoot VSS backup and

restore operations, mailbox restore errors, and VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems.

Troubleshooting VSS backup and restore operations

If you encounter a problem during VSS backup and restore processing, attempt to reproduce the problem in your environment.

Before you begin

If a VSS backup fails, verify that sufficient disk space is available to store the snapshot.

About this task

Procedure

1. Try the operation that failed again.
2. Restart the IBM Spectrum Protect services, including the IBM Spectrum Protect Client Acceptor and the IBM Spectrum Protect Remote Client Agent.
3. If the problem still exists, close other applications, especially those applications, for example antivirus applications, that interact with Exchange Server, SQL Server, or file systems. Retry the operation that failed.
4. If the problem persists, look for information in the event logs: `tdpexc.log`, `tdpsql.log`, and `baclient\dsmerror.log`.
You can also review the messages in the Windows System and Application Event Log. Log entries might exist to help you identify the VSS event that triggers the issue.
5. If you do not find a resolution to the problem in the log files, complete the following steps:
 - a) Restart the Exchange or SQL Server or the computer.
 - b) Run the operation that failed.

Failovers from VSS instant restore processing to VSS fast restore processing

If an error occurs early in a VSS instant restore operation, the error might cause the system to fail over to VSS fast restore processing. However, if an error occurs later in the instant restore operation, instant restore processing might fail without failing over to fast restore processing.

About this task

Errors in VSS instant restore operations might occur, for example, if the volume where the restored database is stored is used by another process.

Procedure

- Check the error message in the `dsmerror.log` file.

Troubleshooting file system and custom application VSS restores from IBM Spectrum Protect server

File system and custom application VSS restores from IBM Spectrum Protect server are volume image-level restore operations. This type of restore operation might cause the shadow copies, which are created with a system provider for the volume that is being restored, to become invalid and be deleted.

About this task

This issue occurs when the shadow storage for the volume is located within the volume. During the volume image-level restore operation, the shadow storage data is overwritten and the shadow copies are invalidated.

Procedure

- Allocate the shadow storage on a different volume.

For example, with the **vssadmin** tool, use the **Add ShadowStorage** command:

```
vssadmin Add ShadowStorage /For=D: /On=F: /MaxSize=your size
```

Troubleshooting issues with SQL Server tail-log backups

A database restore operation might fail if transaction log records in the *tail of the log* are not backed up.

About this task

During the restore operation, you might see the following error message:

```
Failed - An exception occurred while executing a Transact-SQL statement or batch.  
The tail-log backup of the dbName database has not been backed up.  
Use BACKUP LOG WITH NORECOVERY to backup the log if it contains work you do not want to lose.  
Use the WITH REPLACE or WITH STOPAT clause of the RESTORE statement to overwrite the contents of the log.
```

```
RESTORE DATABASE is terminating abnormally.  
Changed database context to 'master'. (HRESULT:0x80131501)
```

To resolve the error, complete the tail-log backup.

Procedure

1. On the **Protect** tab of the SQL instance, click **Show Backup Options** and set the Back Up Tail-log option to **True**.
2. On the **Actions** pane, select **Log Backup to IBM Spectrum Protect**.

Troubleshooting VSS offline restore of a master database

Microsoft SQL Server only supports offline VSS restores of the master database. IBM Spectrum Protect Snapshot for SQL Server does not support offline restore operations. Therefore, you cannot use IBM Spectrum Protect Snapshot for SQL Server to restore the master database.

Procedure

1. Ensure that the SQL Server is online.
2. Restore the master database to a new database in Microsoft Management Console (MMC), or at the command line. For example: Enter the **tdpsqlc** command with the **/recovery=no** option.
3. After the restore operation is complete, verify that all data files are restored successfully.
4. Stop the SQL Server instance, and rename all data files of the master database.
5. Copy all data files from the new master_restore database to the location of the master database. Verify that all data files are copied.
6. Start the SQL Server instance and verify that the master database is restored successfully.

Related tasks

Restoring the master database

A damaged master database can prevent the SQL Server from starting and cause other errors. To protect your data if the master database is damaged, you must routinely complete a full database backup of the master database (msdb).

Troubleshooting VSS limitations with IBM SAN Volume Controller and IBM Storwize V7000

When you run an IBM Spectrum Protect Snapshot for SQL Server VSS backup (non-offloaded) to an IBM Spectrum Protect server, the IBM SAN Volume Controller or IBM Storwize V7000 LUNs can sometimes remain mapped to the Windows host even though the backup is complete.

Procedure

- Use a backup destination other than IBM Spectrum Protect server (BOTH or LOCAL).

Results

When you run two IBM Spectrum Protect Snapshot for SQL Server VSS backups and if the volumes are large, or the background copy rate is set to a low number, or both conditions occur, the second VSS backup might be presented to be in a hang state. Typically, the Exchange Server data is on IBM SAN Volume Controller or IBM Storwize V7000 disks. However, the second backup is waiting for the IBM SAN Volume Controller or IBM Storwize V7000 background copy of the first backup to complete before proceeding. IBM SAN Volume Controller or IBM Storwize V7000 does not allow two background copies of the same volume to occur at the same time. You might not know that the second backup is waiting for the first background copy to complete.

You might also see timeout errors if the previous IBM SAN Volume Controller or IBM Storwize V7000 background copy takes too long.

What to do next

To resolve timeout issues, schedule VSS backups so that enough time elapses between backups, or increase the copy rate of the IBM SAN Volume Controller or IBM Storwize V7000 background copy.

Troubleshooting mailbox restore errors

If you encounter a mailbox restore error, determine whether the problem is reproducible on other Exchange Servers.

About this task

Mailbox restore errors that you might encounter include MAPI connection issues to the mailbox, insufficient role-based access control (RBAC) permissions to complete the restore operation, or issues with the Mailbox Restore Browser feature.

Related tasks

[Restoring mailbox data](#)

IBM Spectrum Protect Snapshot backs up mailbox data at the database level, and also restores individual mailbox items from the database backup.

Troubleshooting insufficient RBAC roles and permissions

For the following mailbox restore errors, ensure that the RBAC roles and management role scope are set on the Exchange objects for the Exchange user.

Procedure

1. If a mailbox fails to open and the error message indicates a missing RBAC permission, ensure that the user who is logged on to the mailbox has the required RBAC roles, and the management scope for those roles includes the database that contains the mailbox. Then, open the mailbox again.
2. If a mailbox restore operation fails and the error message indicates a missing RBAC permission, ensure that the user who is logged on to the mailbox has the required RBAC roles, and the management scope for those roles includes the source and target databases. Then, restart the restore operation.

Related concepts

[Security requirements for backup and restore operations](#)

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

Troubleshooting mailbox permissions, authentication methods, and registry key settings in a Microsoft Exchange 2013 environment

To resolve mailbox restore errors in an Exchange Server 2013 environment, ensure that the Exchange Server mailbox permissions, authentication methods, registry key settings, and the Client Access Server (CAS) role are configured correctly.

Procedure

1. Grant full access permission to the user who is logged on to the target mailbox.
When the administrator mailbox is used, Exchange Server 2013 usually blocks full access permission for the administrator by default.
2. To restore an Exchange 2013 public folder mailbox, ensure that the Exchange user has the Public Folders management role.
3. Log on to an Exchange Server 2013 mailbox as the Exchange Server administrator and ensure that sufficient storage space is available in the administrator mailbox.
4. Ensure that you can access the mailbox that you logged on to and the target mailbox in either Microsoft Outlook or Outlook Web Access.
5. Specify an Exchange Server 2013 CAS by setting the **CLIENTACCESSServer=servername** parameter. If you are using a load balancer, set the **CLIENTACCESSServer** parameter to point to the CAS instead of the load balancer.
6. Open the administrator mailbox and the target mailbox. On the **Actions** pane in the **Mailbox Restore Browser** interface, click **Open Exchange Mailbox**.
7. Verify that the MAPI registry key, `RpcHttpProxyMap_TSM`, is correct to enable IBM Spectrum Protect Snapshot to connect to the Exchange Server. Use one of the following methods:
 - Check the registry key that is in the `HKEY_CURRENT_USER\Software\Microsoft\Windows NT\Current Version\Windows Messaging Subsystem` directory. Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment. For example, you might specify HTTPS as the authentication method if RPC-over-HTTPS connections are enabled for the Exchange Server that is hosting the MAPI profile. Otherwise, you might use HTTP authentication for RPC-over-HTTP connections.
 - Use the **MAPI Settings** property page in Microsoft Management Console (MMC) to ensure that the MAPI registry key is correct. Change the registry key values to reflect the correct domain, endpoint, and Remote Procedure Call (RPC) authentication methods for your environment.

By default, the following registry key format is used.

```
Domain=Proxy Server,RpcHttpAuthenticationMethod,  
RpcAuthenticationMethod,IgnoreSslCert
```

Where the

- *Domain* value is the domain suffix of the personalized server ID, for example, `companyname.local`. Specify any domain or a substring of a domain, or the asterisk (*) and question mark (?) wildcard characters, for example, `*.companyname.local`.
- *Proxy Server* value is the RPC proxy server that has the Client Access Server (CAS) role. Specify the fully qualified domain name (FQDN) of the RPC proxy server. Precede the FQDN by `http://` for an HTTP connection, or `https://` for an HTTPS connection. For example, `https://exchange.companyname.com`
- *RpcHttpAuthenticationMethod* value is the method that is used to authenticate RPC-over-HTTP connections. Specify `NTLM`, `Basic`, `Negotiate`, or `WinNT`.
- *RpcAuthenticationMethod* value is the method that is used to authenticate RPC-over-TCP connections. Specify `NTLM`, `Negotiate`, `WinNT`, `Anonymous`, or `None`.

- *IgnoreSslCert* value indicates whether the Exchange Server validates SSL certificates. For the Exchange Server to ignore invalid certificates, specify `False`.

The default registry key looks like the following example:

```
contoso.com=https://mail.contoso.com,ntlm,ntlm,false
```

Related tasks

[Ensuring successful MAPI connections](#)

If you use Exchange Server 2013, use the **MAPI Settings** property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables IBM Spectrum Protect Snapshot to connect to the Exchange Server.

Troubleshooting MAPI connection issues

Procedure

- To diagnose MAPI-to-mailbox connection issues, enter the **TDPMAPI TESTMAPI** command with these parameters:

/MAILBOXALIAS

Exchange Server 2013: This parameter is the alias name for the mailbox that you are logged on to. The parameter refers to the email alias for the user and is the portion of the email address before the @ symbol. Run this command for the mailbox to be restored and the mailbox that you are logged on to.

Exchange Server 2016 or later: This parameter is the SMTP address for the logged on user's mailbox. The value can be shown by using Exchange cmdlet **Get-Mailbox <mailbox_name> | Select PrimarySmtpAddress**

/EXCSERVER

Exchange Server 2013: This parameter is the name of the mailbox endpoint of the user who is logged in. Use the Exchange Powershell command, **whoami | Get-Mailbox | fl ExchangeGUID**, to determine the value. You must specify this parameter for Exchange Server 2013.

Exchange Server 2016 or later: This parameter is the alias name for the mailbox that you are logged on to.

/TRACEFILE

This parameter is the file name that is used to store the output from tracing operations. By default, tracing is turned off. You can qualify the file name by specifying a drive and a full directory path. You must have write permissions for the user that runs the command.

Related tasks

[Ensuring successful MAPI connections](#)

If you use Exchange Server 2013, use the **MAPI Settings** property page to verify that the user mailbox is online. You can also view and update the MAPI registry key that enables IBM Spectrum Protect Snapshot to connect to the Exchange Server.

Troubleshooting a MAPI error that prevents multiple mailboxes restoring in a Microsoft Exchange 2013 environment

When you restore multiple mailboxes on a server that is running Exchange Server 2013, the mailbox restore operation might partially fail and report a MAPI error.

About this task

In Exchange Server 2013, Client Throttling Policy (the **RcaMaxConcurrency** parameter), specifies how many concurrent connections you can maintain at one time. If you attempt to make more concurrent requests than the **RcaMaxConcurrency** parameter allows, the new connection attempt fails. However, the existing connections remain valid.

Procedure

- Increase the **RcaMaxConcurrency** value for the logon user mailbox.
For more information about this setting, see Microsoft documentation: [Exchange 2013 Client Throttling \(http://technet.microsoft.com/en-us/library/bb232205\(v=exchg.150\).aspx\)](http://technet.microsoft.com/en-us/library/bb232205(v=exchg.150).aspx)

Troubleshooting issues with the Mailbox Restore Browser interface on remote systems

An error can occur when many mailboxes are queried on a remote system, which causes an out-of-memory exception.

About this task

If you restore mailboxes on the remote system, the list of mailboxes might not be displayed in the Source mailbox navigation tree of MMC. You might see the following message:

```
Error: Processing data for a remote command failed with the following error message:
```

```
The WSMAN provider host process did not return a proper response. A provider
```

```
in the host process may have behaved improperly. For more information, see the
```

```
about_Remote_Troubleshooting Help topic.OperationStopped: (<Machine_Name>:String)
```

```
[],PSRemotingTransportExceptionJobFailure
```

Procedure

- To resolve the out-of-memory exception, increase the default memory value for the remote Powershell session.

For example, to increase the maximum of memory that is allocated per shell to 4 GB, enter the following cmdlets at the Microsoft Windows PowerShell command line.

```
Set-Item WSMAN:\localhost\Shell\MaxMemoryPerShellMB 4096
```

```
Set-Item WSMAN:\localhost\Plugin\Microsoft.PowerShell\Quotas\MaxMemoryPerShellMB  
4096
```

```
Restart-Service winrm
```

Troubleshooting an SMTP restore issue that occurs when you restore email with large attachments in the Mailbox Restore Browser interface

If you restore an email with an attachment that is larger than 3 MB to an SMTP server, a Microsoft fix is required.

About this task

You might see the following error message:

```
QFD: System.Net.Mail - SmtplibClient class throws exceptions if file attachment
```

```
is over 3 MB
```

Procedure

- Resolve the issue by applying the fix that is available at this web page: [Microsoft Connect Visual Studio and .NET Framework Downloads \(http://support.microsoft.com/kb/2183292\)](http://support.microsoft.com/kb/2183292)

Troubleshooting a limitation with deleted mailbox history in the Mailbox Restore Browser interface

IBM Spectrum Protect Snapshot does not record the time when mailboxes are deleted.

About this task

After a mailbox is deleted, the **Available Database Backups** list in the Mailbox Restore Browser continues to list database backups that contained the mailbox before its deletion.

Procedure

- Ensure that the backup version that you select to restore from the **Available Database Backups** list contains the mailbox before it was deleted.

Deleting mailbox history information

Mailbox history includes only the mailboxes from databases that are backed up. If you back up mailbox history with a version of IBM Spectrum Protect Snapshot earlier than version 4.1, you can manually delete the old mailbox history.

About this task

IBM Spectrum Protect Snapshot for Microsoft Exchange Server backs up a new set of mailbox history data. With the new mailbox history data, you can experience better performance when you back up mailbox history. It is also easier to find the mailbox when you restore a mailbox. Additionally, when you retrieve mailbox history, the mailbox names can be displayed in multiple languages.

Deleting the old mailbox history is not required. If you delete the old mailbox history data, you lose the location history information for the deleted and moved mailboxes in the backup copies that earlier versions of IBM Spectrum Protect Snapshot for Microsoft Exchange Server created.

Even if a mailbox user is deleted from Active Directory and backups that contained that mailbox are expired, IBM Spectrum Protect Snapshot retains the mailbox history information indefinitely on the IBM Spectrum Protect server. Therefore, you can still see the mailbox history information for deleted mailboxes within the restore search views even though the associated backups might be expired. The mailbox restore list, which is populated from the mailbox history, is not intended to be an all-inclusive list of mailboxes that can be restored. It is made available for ease of use.

Procedure

1. Enter the following command to save the mailbox history to a file:

```
tdpexcc q tsm /showmailboxinfo > E:\MyMailboxHistory.txt
```

Keep this file for reference. You can use the backup copy when you need location information for the deleted and moved mailboxes

2. If you need to restore a mailbox from the old backup copies, and the mailbox location changes before you delete the mailbox history, use the **/MAILBOXORIGLOCATION** parameter to restore the mailbox. After the old backup copies expire, mailbox history works without you having to specify the **/MAILBOXORIGLOCATION** parameter.
3. Complete the following steps to delete the old mailbox history from the IBM Spectrum Protect server.
 - a) Start the IBM Spectrum Protect command-line administrative interface, `dsmadm . exe`.
 - b) Log on to the IBM Spectrum Protect server.
 - c) Enter the following command to query the filespace name:

```
Query Filespace node_name file_space_name
```

The format of the filespace name for mailbox history is *DomainName\MAILBOXINFO*. For example, the following command queries the filespace for the mailbox history for the *CXCLAB_EXC* node. The

node_name is the **DAGNODE** name, or the Exchange Server node name when the **DAGNODE** is not being used.

```
tsm: FCM>QUERY FILESPACE CXCLAB_EXC *MAILBOXINFO
```

The following results are displayed:

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity	Pct Util
CXCLAB_EXC	cxcserver.com\MAILBOXINFO	52	TDP MSE-xchg	API:ExcData	No	0 KB	0.0

4. Enter the following command to delete the filesystem for the old mailbox history while bearing in mind that all previous backups might be deleted if you do not enter the command correctly.

```
DELEte Filespace node_name file_space_name\MAILBOXINFO
```

For example, the following command deletes the filesystem for the mailbox history for the *CXCLAB_EXC* node:

```
tsm: FCM>DELETE FILESPACE CXCLAB_EXC cxcserver.com\MAILBOXINFO
```

Related concepts

[Software requirements for mailbox restore operations](#)

When you restore mailboxes and mailbox data, you can choose where to restore the mail and how to restore the mail. You can restore mailbox data from the GUI or command-line interface.

Troubleshooting configuration errors in a failover clustered environment

If you encounter errors when you configure a failover clustered environment, determine whether the options in the backup-archive client and application-specific `dsm.opt` files are specified correctly.

Procedure

1. When you are configuring IBM Spectrum Protect Snapshot as a stand-alone configuration, verify that the same path to the `VSSALTSTAGINGDIR` directory is specified in the backup-archive client options file, `baclient\dsm.opt`, and in the IBM Spectrum Protect Snapshot for SQL Server options file, `tdpsql\dsm.opt`.
2. Verify that the `VSSALTSTAGINGDIR` path in the `dsm.opt` files points to a directory on a shared disk or cluster shared volume that all cluster nodes can access.
3. In the Data Protection and IBM Spectrum Protect Snapshot `dsm.opt` files, verify that the option for the **CLUSTERNODE** parameter is set to yes.

When you work in a cluster environment, from the command-line interface, the Volume GUID is not displayed for volumes that are clustered disks. The clustered disk is displayed.

Related tasks

[Configuring IBM Spectrum Protect Snapshot for SQL Server clustered environments](#)

Depending on the applications that are installed, you can configure IBM Spectrum Protect Snapshot to operate in clustered environments with SQL Server, or file system workloads.

Troubleshooting VSS and SAN Volume Controller, Storwize V7000, or DS8000

If you experience VSS and SAN Volume Controller, Storwize V7000, or DS8000 problems, investigate the configuration in your environment.

Procedure

1. Verify connectivity to the CIMOM (Common Information Model Object Manager) as follows:

- a) Refer to your SAN Volume Controller, Storwize V7000, or DS8000 documentation.
- b) Run the **IBMVCFG LIST** command. The default location is %Program Files%\IBM\Hardware Provider for VSS-VDS.
- c) Issue the **IBMVCFG SHOWCFG** command to view the provider configuration information.
- d) Check that the CIMOM is properly configured. Run `verifyconfig.bat -u username -p password` on the Master Console.
- e) Check the user name and password.

If the problem is with the truststore, follow the procedure in the documentation to generate a new truststore.

2. Verify CIMOM operational issues as follows:

- a) If your backup or restore operation fails, check the `IBMVSS.log` file.

If the backup or restore failure is from a CIMOM failure, the log displays output similar to the following example:

```
Wed Jan 13 17:34:34.793 - Calling AttachReplicas
Wed Jan 13 17:34:35.702 - AttachReplicas: 909ms
Wed Jan 13 17:34:35.702 - returnValue: 34561
Wed Jan 13 17:34:35.718 - AttachReplicas returned: 34561
java.util.MissingResourceException: Can't find resource for
bundle java.util.PropertyResourceBundle, key 1793
at java.util.ResourceBundle.getObject(ResourceBundle.java:329)
at java.util.ResourceBundle.getString(ResourceBundle.java:289)
at com.ibm.cim.CIMException.<init>(CIMException.java:472)
at ESSService.executeFlashCopy(ESSService.java:3168)
Wed Jan 13 17:34:35.779 - IBMVSS: AbortSnapshots
```

A return value of 0 means that the backup or restore operation is successful.

- b) To determine why a backup or restore operation failed, review the log files.

Tip: If VSS backups fail, issue the **IBMVCFG LIST FREE** command verify that sufficient free volumes are available in the VSS_FREE volume group to store the snapshot.

3. If the failure seems to be for a different reason than a CIMOM failure, verify your host configuration. Run the latest support levels of the software for SAN Volume Controller, Storwize V7000, or DS8000.
4. If you are unable to resolve these problems, provide the following information to IBM Support:

- Information that is listed in the IBM Spectrum Protect diagnostic information section
- HBA type, firmware, and driver levels
- SDD version
- SAN Volume Controller microcode version (if applicable)
- DS8000 microcode version (if applicable)
- Storwize V7000 microcode version (if applicable)
- SAN Volume Controller or Storwize V7000 Master Console version (if applicable)
- For DS8000, the CIM Agent version (if applicable)
- `IBMVSS.log`
- `IBMVDS.log`
- Application Event Log
- System Event Log
- CIMOM logs if the problem seems to be related to CIMOM. Run `CollectLogs.bat` and send the file that is created (`CollectedLogs.zip`) to IBM Support.

The default location for SAN Volume Controller or Storwize V7000 is `C:\Program Files\IBM\svconconsole\support`, and the default location for DS8000 is `C:\Program Files\IBM\cimagent`.

Related concepts

IBM Spectrum Protect Snapshot with IBM SAN Volume Controller and IBM Storwize V7000

The way in which you configure the VSS hardware provider and how you configure your IBM SAN Volume Controller and IBM Storwize V7000 device dictates the characteristics of your VSS snapshots. VSS snapshots are implemented as flashcopies on an IBM SAN Volume Controller and IBM Storwize V7000 device. FlashCopys consist of a source volume and a target volume.

Resolving problems with IBM Support

Contact IBM Support for further assistance if you have a problem that you are unable to solve by applying maintenance fixes, reproducing the issue, or reviewing the information in previous topics. IBM Support might request to see some or all of the trace and log files that are related to a problem that you report.

About this task

Go to the [IBM Support](#) page for IBM Spectrum Protect Snapshot and log in to access support information for the product.

You might be asked to set a trace on the Data Protection client that uses VSS technology, and then collect the log. IBM Support uses the information that is captured in the log file to trace a problem to its source or to determine why an error occurred.

Viewing trace and log files

IBM Spectrum Protect Snapshot uses several components. Each component is in its own directory along with its respective troubleshooting files. By using the **Trace and Log Files** view, you can easily view these files in a central location.

About this task

You can collect trace and log files in the **Diagnostics** property page for a workload.

These diagnostics property pages can control the tracing settings for all related components such as the workload, the IBM Spectrum Protect API, the Client Agent service, and Microsoft Management Console (MMC).

The following diagnostic modes are available:

Normal

Use for SQL legacy backup operations. Using this mode results in a small sized trace file.

Complete (default)

Using this mode results in a large sized trace file.

Custom

Use when full control over trace flags must be set

Procedure

1. When you encounter a problem in MMC, create trace files by using the **Diagnostics** property page.

a) Click **Properties > Diagnostics**, and click **Begin**.

You can set the following items:

- You can click **Screen shot** to open the **Diagnostics screen shot** tool window. When you want to create a screen capture of any open windows, click **Add New Screenshot**. The name of the screen capture is added to the list of items on the **Diagnostics** property page. Close the **Diagnostics screen shot** when you finish taking screen captures.
- For SQL workload instances, enter a database name in the **SQL Database** field, and click **Add Database Information**. Repeat this step as needed. This step is useful if one database can be backed up and another cannot. By providing the details for both databases, it helps identify differences in database properties.

b) Close the property page and reproduce the problem.

c) Open the **Diagnostics** property page and click **Stop**.

Clicking the **Diagnostics** button is the preferred method for gathering information to send to your service representative. This method gathers all the information that is needed. Even if a problem occurs only on the command-line interface, command, you can always gather information by using the **Automate** tab.

The log files are displayed in the **Trace and Log Files** view.

2. Click the trace or log file that you want to view.

The contents of the file are displayed in the results pane. The following files are examples of the files that you can view, including default log and trace files:

Examples of IBM Spectrum Protect Snapshot default log and trace files:

- Installation directory: C:\Program Files\Tivoli\FlashCopyManager
- dserror.log
- Log file for custom applications and file systems workloads: fcm.log
- TraceFm.trc
- TraceUx.trc
- TraceManagedCapacityHistory.trc
- TraceSchedLaunch.trc
- VssProvisioning.log
- TraceFileFS.trc
- TraceFileExc.trc
- TraceFileSql.trc

If the fcm.log is defined in a path other than the default C:\Program Files\Tivoli\FlashCopyManager\fcml.log, the reports do not include the following information for scheduled backup and restore operations:

- Task completion
- Type of data protection activity
- Amount of data protection activity

The charts and reports display only information that is present in the default log file fcm.log.

Examples of trace logs and scripts to quiesce custom applications:

- Default directory: %ALLUSERSPROFILE%\Application Data\Tivoli\FlashCopyManager\custom-application

where *custom-application* can be various applications. For example, an IBM Domino® mail server, or MySQL or Oracle server.

Examples of VSS Requestor default log and trace files:

- Installation directory: C:\Program Files\Tivoli\TSM\baclient
- dserror.log

Examples of IBM VSS provider for SAN Volume Controller, Storwize V7000, and DS8000 log files

- IBMVDS.log
- IBMVss.log

Gathering trace and log files for remote systems

Collecting diagnostic data for a remote system, by using IBM Spectrum Protect Snapshot, is different to collecting data for a local system. You can update the **Diagnostics** property page to collect the correct log and trace files for remote systems.

Before you begin

On the local system, verify the following system requirements:

- Windows 7, Windows 8, Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2 is installed
- PowerShell version 3.0 or later is installed, if you are running Windows 7, Windows 8, Windows 2008, or Windows 2008 R2. With Windows 2012, PowerShell version 4.0 is installed by default.
- IBM Spectrum Protect Snapshot version 4.1.4 is installed

On the remote system, verify the following system requirements:

- Windows 2008, Windows 2008 R2, Windows 2012, Windows 2012 R2, or a later version is installed
- Windows PowerShell version 3.0 or later is installed, if you are running Windows 2008, or Windows 2008 R2. With Windows 2012, PowerShell version 4.0 is installed by default.
- IBM Spectrum Protect Snapshot version 4.1.4 is installed
- The required workload is configured.

Procedure

1. In the Actions pane, click **Properties > Diagnostics**, and select the mode that you require as follows.
 - For a smaller trace file, select **Normal**.
 - For a larger trace file, select **Complete**.
 - For full control over the trace flags that are set, select **Custom**.
2. Click **Begin**.
3. Click **OK** to close the window.
4. Reproduce the issue that you are seeing on the remote server. For example, back up or restore data on the remote Exchange Server.
5. Open the **Diagnostics** property page and click **Screenshot**.

Clicking the **Diagnostics** button is the preferred method for gathering information to send to your service representative. This method gathers all the information that is needed.
6. In the **Diagnostic Screen Shot Tool** window, click **Add New Screen Shot**. An image is displayed.

When you use the **Diagnostic Screen Shot Tool** on the remote system, the screen capture files are on the local system.
7. Close the **Diagnostic Screen Shot Tool** window.
8. Click **End**.

Results

Log, configuration, and trace files are detected and displayed, such as those files in the following example. The diagnostic log files are on the remote system.

- Microsoft Management Console (MMC): TraceFm.trc and TraceUx.trc
- Data Protection: TraceFileFs.trc, TraceFileSql.trc, TraceFileExc.trc
- Agent: TraceFileAgent.trc
- API: TraceFileFsAPI.trc
- Other: Hardware provider logs, System information

The files and traces are stored in the following folder on the remote system: C:\Program Files\Tivoli\FlashCopyManager\ProblemDetermination. Use the **Copy** function to copy the files locally.

If you enable command-line interface tracing, the command-line interface generates trace files for the local and remote systems. On the local system you can view the file that you specified. In addition, on the local and remote systems, a trace file is also created. This file has the same name as the file stored on the local system and the file name concludes with the following suffix appended to the file type extension: *_remote*

For example, on the local system, the following files are created after you enable command-line interface tracing:

- *filename.trc*
- *filename.trc_remote*

On the remote system, the following file is created after you enable command-line interface tracing, *filename.trc_remote*.

Related tasks

Mounting VSS snapshots to remote servers

You can mount VSS snapshots to remote servers that other users can access. When you remotely mount a snapshot, you must use hardware that includes a feature for creating transportable snapshots. If you use a mix of hardware, all hardware must include a feature for creating transportable snapshots.

Gathering information about Exchange or SQL Server with VSS before you call IBM

The Data Protection client depends on the operating system and the Exchange or SQL Server application. Collecting all the necessary information about the environment can significantly assist Support in determining the source of problem.

Procedure

- Gather as much of the following information as possible before you contact IBM Support:
 - The exact level of the Windows operating system, including all service packs and test fixes that were applied.
 - The exact level of the Exchange Server or SQL Server, including all service packs and test fixes that were applied.
 - The exact level of IBM Spectrum Protect Snapshot for Microsoft Exchange Server, or IBM Spectrum Protect Snapshot for Microsoft SQL Server, with Volume Shadow Copy Service (VSS) backup and restore support.
 - The exact level of the IBM Spectrum Protect API.
 - The exact level of the IBM Spectrum Protect server.
 - The exact level of the IBM Spectrum Protect backup-archive client.
 - The exact level of the IBM Spectrum Protect storage agent (if LAN-free environment).
 - The IBM Spectrum Protect server and operating system level.
 - The output from the IBM Spectrum Protect server **QUERY SYSTEM** command.
 - The output from the IBM Spectrum Protect Snapshot for Microsoft Exchange Server **TDPEXCC QUERY EXCHANGE** command.
 - The device type (and connectivity path) of the Exchange Server databases and logs.
 - (SAN only) The specific hardware that is being used. For example: HBA, driver levels, microcode levels, SAN Volume Controller or Storwize V7000 levels, DS8000 hardware details.
 - Permissions and the name of the user ID being used to run backup and restore operations.
 - The name and version of antivirus software.
 - (SAN only) The VSS hardware provider level.

- The VSS hardware provider log files. See the documentation of the specific VSS hardware provider on how to enable tracing and collect the trace log files.
- (SAN only) The IBM CIM agent level for DS8000, SAN Volume Controller, or Storwize V7000.
- A list of vendor-acquired Exchange applications that are running on the system.
- A list of other applications that are running on the system.
- A list of the steps that are needed to re-create the problem (if the problem can be re-created).
- If the problem cannot be re-created, list the steps that caused the problem.
- Is IBM Spectrum Protect Snapshot for Microsoft SQL Server running in a Microsoft Failover Clustering environment?
- Does the problem occur on other Exchange or SQL servers?

Viewing system information

You can view and edit scripts that provide information about system components such as Windows-related services for IBM Spectrum Protect Snapshot, Windows event log entries, and Volume Shadow Copy Service (VSS) information.

About this task

The **System Information** view is extensible. You can take advantage of this flexibility to add and share customize scripts.

Procedure

1. Open the **System Information** view as follows:

- a) Click **Diagnostics** in the results pane of the welcome page.
- b) Double-click **System Information** in the results pane.

A list of scripts is displayed in the results pane of the **System Information** view. The types of scripts that are displayed are PowerShell scripts, Windows Management Instrumentation scripts, and IBM Spectrum Protect scripts.

2. Add, update, or delete your scripts, as follows:

Action	Steps
Add your own scripts	<ol style="list-style-type: none"> a. Click New in the Actions pane. b. If you want to copy your scripts directly to the ProgramFiles\Tivoli\FlashCopyManager\Scripts directory, make sure that your scripts follow these extension requirements: <ul style="list-style-type: none"> • PowerShell scripts: <i>filename.ps1</i> • Windows Management Instrumentation (WMI) scripts: <i>filename.wmi</i> • IBM Spectrum Protect scripts: <i>filename.tsm</i> <p>IBM Spectrum Protect Snapshot uses the file type extension to determine how to run the script.</p>
View or edit an existing script	<ol style="list-style-type: none"> a. From the list of script files in the results pane, select the name of a script that you want to view or edit. <p>Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to process.</p> b. To open the script file for viewing or editing, click Command Editor in the Actions pane. c. View or edit the script.

Action	Steps
	d. Click OK to save your changes, or click Cancel to exit the System Information Command Editor without saving any changes.
Delete a script	<p>a. From the list of script files in the results pane, select the name of a script that you want to delete.</p> <p>Tip: The name of the script is displayed in the Actions pane. Click the name of the script in the Actions pane to reveal or hide a list of actions to process.</p> <p>b. Click Delete in the Actions pane.</p>

Emailing files to IBM Support

You can send diagnostic information to IBM Support.

Before you begin

About this task

The email support files feature collects all detected configuration, option, system information, trace, and log files. It also collects information about services, operating systems, and application versions. These files are compressed and then attached in an email.

Procedure

1. Start Microsoft Management Console (MMC).
2. Click **Diagnostics** in the results pane of the welcome page.
3. In the **Actions** pane, click **E-Mail Support files**.
4. Enter the information in the various fields and click **Done**.

The information is sent to the designated support personnel and the dialog closes.

Online IBM support

Multiple online support resources are available for your reference.

The following list identifies where you can find information online:

- Tivoli Storage Manager wiki (<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager>).
- Storage Management community on Service Management Connect (<https://www.ibm.com/developerworks/servicemanagement/sm/index.html>).
- IBM Spectrum Protect Snapshot (www.ibm.com/software/products/en/spectrum-protect-snapshot). Enter the search term to narrow the search criteria for your support requirements. Examples of search terms that you might use include an authorized program analysis report (APAR) number, release level, or operating system.
- IBM Spectrum Protect for Mail (www.ibm.com/software/products/en/spectrum-protect-for-mail). Enter the search term to narrow the search criteria for your support requirements. Examples of search terms that you might use include an authorized program analysis report (APAR) number, release level, or operating system.
- IBM Spectrum Protect for Databases (<http://www.ibm.com/software/products/en/spectrum-protect-for-databases>). Enter the search term to narrow the search criteria for your support requirements. Examples of search terms that you might use include an authorized program analysis report (APAR) number, release level, or operating system.

Chapter 8. Reference

Reference topics provide information related to IBM Spectrum Protect Snapshot. Topics include information about the backup and restore commands that you can enter at the command-line interface as an alternative to using Microsoft Management Console (MMC).

Support for Microsoft Exchange 2016 and later versions

With IBM Spectrum Protect Snapshot for Microsoft Exchange Server version 4.1.4.2, features that support Microsoft Exchange 2016 were added and you can now protect and manage your Microsoft Exchange 2016 environment.

Mailbox filter options

When you restore an individual mailbox, you can use mailbox filters to identify individual messages to restore. With Microsoft Exchange 2016 or later versions, the **Folder Name** filter option is supported.

For example, to restore a folder that is named “folder A” in the mailbox “MailboxA”, run the following command:

```
tdpexcc restoremailbox "MailboxA" /MailboxFilter="folder, folderA"
```

Individual mailbox restore options

You can restore individual mailbox items from database backups. The following table describes the differences between the mailbox restore features that are supported with Microsoft Exchange 2013, and the features that are supported with Microsoft Exchange 2016 and later versions.

Feature	Description	Exchange 2013	Exchange 2016 or later
Mailbox restore	Mailbox restore browser	Only supports Non-Unicode PST file.	Only supports Unicode PST file.
	Non-Unicode mailbox restore	The Restore Mailbox to non-Unicode PST file option is available for selection in the Actions pane.	Not supported with Microsoft Exchange 2016 or later versions. Note: The mailbox is automatically restored to a Unicode PST file.

Temporary mailbox folder cleanup

When the mailbox is successfully restored, with IBM Spectrum Protect Snapshot for Microsoft Exchange Server version 4.1.4.2, the temporary mailbox folder that was created during the restore operation can be deleted automatically.

Note: To enable automatic temporary folder deletion with Microsoft Exchange 2016, you must log on as an Exchange Server administrator and ensure that the **ApplicationImpersonation** role is assigned to your user. This role is not enabled by default.

Message application programming interface (MAPI) client and collaboration data objects (CDO)

The MAPI/CDO library is not supported with Microsoft Exchange 2016 or later versions. The MAPI/CDO changes table describes the impact of this change in your IBM Spectrum Protect Snapshot for Microsoft Exchange Server solution.

Table 38. MAPI/CDO changes

Feature	Description	Exchange 2013	Exchange 2016 and later versions
MAPI Settings	MAPI Settings property page	The MAPI Settings property page is available under the Protect and Recover Data node on the MMC.	Not supported with Microsoft Exchange 2016 or later.
	MAPI configuration checks that use the configuration wizard	When you use the configuration wizard to configure IBM Spectrum Protect Snapshot for Microsoft Exchange Server on the MMC, the system automatically runs a number of checks to verify that the Microsoft Exchange Server MAPI client and CDO are correctly installed.	When you use the configuration wizard to configure IBM Spectrum Protect Snapshot for Microsoft Exchange Server on the MMC, the system automatically runs a number of checks to verify that the correct version of Microsoft Outlook is installed.

Command-line overview: IBM Spectrum Protect Snapshot for Exchange Server

The name of the IBM Spectrum Protect Snapshot for Exchange Server command-line interface is `tdpexcc.exe`. If you installed the **TDPEXchange** package, or you configured the Exchange Server in Microsoft Management Console (MMC), the program is (by default) in the IBM Spectrum Protect Snapshot for Exchange Server installation directory (`C:\Program Files\Tivoli\tsm\TDPEXchange\`).

Command-line parameter characteristics

The command-line parameters have the following characteristics:

- Positional parameters do not include a leading slash (/) or dash (-) character.
- Optional parameters can display in any order after the required parameters.
- Optional parameters begin with a forward slash (/) or a dash (-) character.
- Minimum abbreviations for keywords are indicated in uppercase text.
- Some keyword parameters require a value.
- For those keyword parameters that require a value, the value is separated from the keyword with an equal sign (=) character.
- If a parameter requires more than one value after the equal sign, the values are separated with commas.
- Each parameter is separated from the others by using spaces.
- If a parameter value includes spaces, the value must be enclosed in double quotation marks.
- A positional parameter can display only once per command invocation.

Command-line interface help

Issue the **tdpexcc ?** or **tdpexcc help** command to display help for the command-line interface. You can see more specific help for commands by entering a command like the following example: **tdpexcc help backup**, where **backup** is an example of a command.

Related tasks

[Protecting Microsoft Exchange Server data](#)

With IBM Spectrum Protect Snapshot for Exchange Server, you can back up and restore Microsoft Exchange Server databases.

Backup command

Use the **backup** command to run Exchange Server backups of databases from the Exchange Server to local shadow volumes managed by IBM Spectrum Protect Snapshot.

You must have local registry rights (for all versions of Exchange Server) to run a IBM Spectrum Protect Snapshot for Exchange Server backup.

Microsoft Exchange Server considers the asterisk (*) wildcard character to be an invalid character when used in database names. Databases that contain the asterisk (*) wildcard character in their name are not backed up. When a full VSS snapshot backup is done, the backup remains active until the backup version is deleted with the delete backup command, or expired by IBM Spectrum Protect Snapshot according to the defined policy. Two different active backups can exist at the same time:

- Full backup, along with any associated incremental backups and differential backups.
- Copy backup, along with any associated incremental backups and differential backups.

When you run Exchange Server backups, the Exchange database file size might increase because of increased database commitments that are triggered by backup operations. This condition is a Microsoft Exchange server standard behavior.

IBM Spectrum Protect Snapshot for Exchange Server supports the following types of VSS backups:

Full

Back up the entire database and transaction logs. If a successful backup is obtained, the Exchange Server deletes the committed log files. In Exchange Server Database Availability Group environments, the log files might not be immediately deleted after a successful full backup.

Incremental

Back up the transaction logs. If a successful backup is obtained, the Exchange Server deletes the committed log files. In Exchange Server Database Availability Group environments, the log files might not be immediately deleted after a successful incremental backup.

Differential

Back up the transaction logs. The transaction logs are not deleted.

Copy

Back up the entire database and transaction logs. The transaction logs are not deleted.

All databases must be mounted at the time of the backup operation. If any database is not mounted, the database is not backed up. In addition, the transaction logs are not truncated.

Backup syntax

Use the **backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Backup positional parameters

Positional parameters immediately follow the **backup** command and precede the optional parameters.

The following positional parameters specify the object to back up:

* | *dbname1, ..., dbnameN*

*

Back up all databases.

dbname

Back up the specified database. Multiple entries are separated by commas. If separated by commas, ensure that there is no space between the comma and the name. If any database name contains blanks, enclose the database name in double quotation marks.

The following positional parameters specify the type of backup to run:

FULL | **COPY** | **INCR**emental | **DIFF**erential

FULL

Back up the entire database, and the transaction logs, and if a successful backup is obtained, truncate the transaction logs.

COPY

Back up the entire database, and the transaction logs, and do NOT truncate the transaction logs.

INCRemental

Back up the transaction logs, and if a successful backup is obtained, truncate the transaction logs.

DIFFerential

Back up the transaction log files, but do not truncate the log files.

Backup optional parameters

Optional parameters follow the **backup** command and positional parameters.

/BACKUPDESTination=**LOCAL**|**TSM**|**BOTH**

When you are backing up data to a local system, set **BACKUPDEST**ination to LOCAL. When you are backing up data to an IBM Spectrum Protect server, set **BACKUPDEST**ination to TSM. To back up data to a local system and an IBM Spectrum Protect server, set the parameter to BOTH.

/CONFIGfile=*configfilename*

Use the **/CONFIG**file parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values to use for a **backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/CONFIG**file parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIG**file parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/DAGNode=*nodename*

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Exchange Server from making too many backups of the same database.

/EXCLudedAGActive

Use the **/EXCL**udedAGActive parameter to exclude the Exchange Server databases from backup if they belong to a Database Availability Group and are an active database copy.

/EXCLUDEDAGPASSive

Use the **/EXCLUDEDAGPASSive** parameter to exclude the Exchange Server databases from backup if they belong to a Database Availability Group and are a passive database copy.

/EXCLUDEDB=db-name1,db-nameN,...

Use the **/EXCLUDEDB** parameter to exclude the specified Exchange Server databases from the backup operation. If the database names are separated by commas, ensure that there are no spaces between the commas and the database names. If any database name contains blanks, enclose the database name in quotation marks. You cannot specify the asterisk (*) wildcard character.

/EXCLUDENONDAGDbs

Use the **/EXCLUDENONDAGDbs** parameter to exclude the Exchange Server databases from backup if they do not belong to a Database Availability Group.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, some days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.

- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MINimumbackupinterval=numberOfMinutes

If you are scheduling the backup of databases in an Exchange Server Database Availability Group, specify the minimum amount of time, in minutes, before a backup of another copy of the same Database Availability Group database can begin. The range is 1 - 9999.

Setting this parameter specifies that only one database copy can be backed up within a time frame. This option prevents all of the members in a Database Availability Group from backing up the database, which would be redundant and invalidate the IBM Spectrum Protect storage management policy.

/MOUNTWait=Yes | No

Use the **/mountwait** parameter to specify whether waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify these options:

Yes

IBM Spectrum Protect Snapshot for Exchange Server waits until all initial volumes of any required removable media are made available to the IBM Spectrum Protect server before it completes the command. This option is the default.

No

IBM Spectrum Protect Snapshot for Exchange Server ends the command (if removable media are required). An error message is displayed.

/OFFLOAD

Specify this option if, after the VSS snapshot is complete, you want to offload the transfer of the data from the IBM Spectrum Protect server to the system specified by the **REMOTEDSMAGENTNODE** parameter. This option is only valid when the **BACKUPDESTination** parameter is set to either TSM or BOTH. The default is to not offload data.

/PREFERDAGPASSive

If you are scheduling the backup of databases in an Exchange Server Database Availability Group, set this parameter to back up a passive database in an Exchange Server Database Availability Group unless no valid passive copy is available. If no valid passive copy is available, the backup is created from the active database copy.

/Quiet

This parameter prevents status information from being displayed. This function does not affect the level of information that is written to the activity log.

/SKIPINTEGRITYCHECK

Use the **/SKIPINTEGRITYCHECK** parameter to specify whether IBM Spectrum Protect Snapshot for Exchange Server bypasses the integrity checking of databases and log files, or automatically runs the integrity checking of databases and log files.

You can specify the following values:

No

Run integrity checking to verify that all database and log files do not contain integrity issues. This option is the default.

Yes

Bypass integrity checking of all database and log files during backup processing.

SKIPDBCHECKIFHEALTHY

Bypass integrity checking of database files only if at least two healthy copies of a database (one active and one passive copy) exist in a Database Availability Group (DAG).

SKIPDBANDLOGCHECKIFHEALTHY

Bypass integrity checking of all database and log files during backup processing only if at least two healthy copies of a database (one active and one passive copy) exist in a DAG.



Attention: If you do not specify a value with the SKIPINTEGRITYCHECK parameter, integrity checking of database and log files is bypassed. If you bypass integrity checking, the backup that is stored on IBM Spectrum Protect server might not be valid, or data loss can occur.

/TSMOPTFile=tsmoptfilename

The **/TSMOPTFile** parameter specifies the IBM Spectrum Protect Snapshot options file to use. Considerations:

- The *tsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *tsmoptfilename* variable contains spaces, enclose the variable in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\dsm.opt"
```

- If you do not specify **/TSMOPTFile**, the default value is *dsm.opt*.
- If you specify **/TSMOPTFile**, but not *tsmoptfilename*, the default is also *dsm.opt*.

/UpdateMailboxInfoOnly

Specify the **/UpdateMailboxInfoOnly** parameter to update only the mailbox history information in Microsoft Exchange Server database backups, for example:

```
tdpexcc backup DB1 full /UpdateMailboxInfoOnly
```

where DB1 is the database name, and full is the type of database backup.

Restriction: This parameter does not back up the Exchange Server database.

Backup example

The following list provides examples of how to use the **backup** command.

To complete a full backup of a database, for example, *DB_G*, the following command can be entered:

```
tdpexcc backup DB_G full
```

To complete a copy backup of a database, for example, *DB_G*, the following command can be entered:

```
tdpexcc backup DB_G copy
```

To complete a differential backup of a database, for example, *DB_G*, the following command can be entered:

```
tdpexcc backup DB_G diff
```

To complete an incremental backup of a database, for example, *DB_G*, the following command can be entered:

```
tdpexcc backup DB_G incr
```

Delete backup command

Use the **delete backup** command to delete a VSS backup of an Exchange Server database.

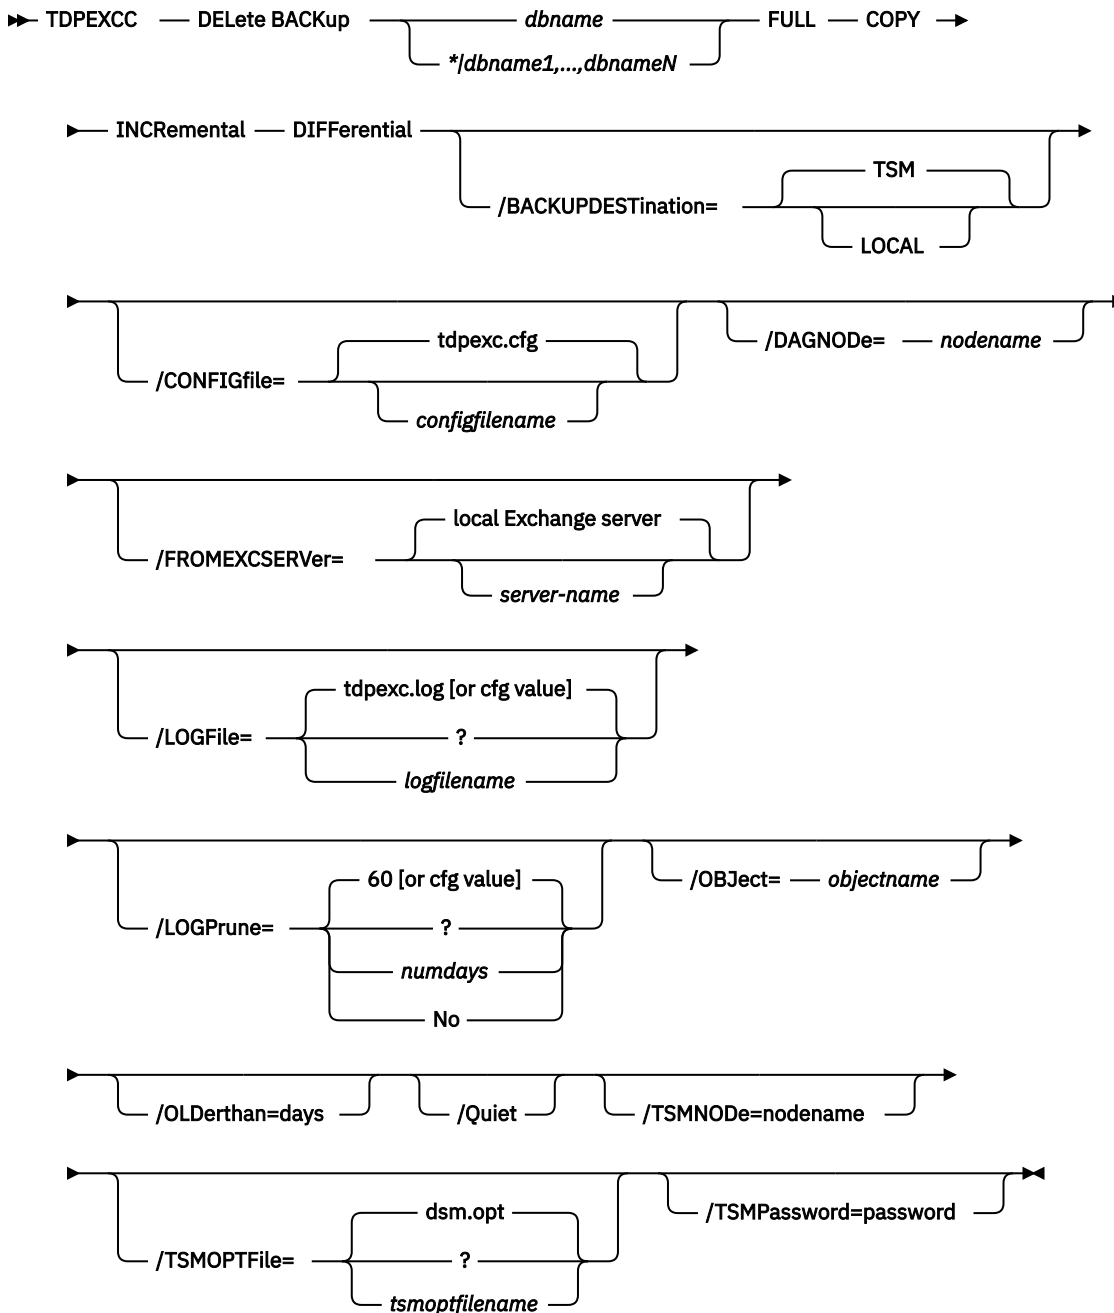
You must have local registry rights (for all versions of Exchange Server) to run a IBM Spectrum Protect Snapshot for Exchange Server delete backup.

- When you run a full VSS snapshot backup, the backup remains active until the backup version is either deleted with the delete backup command, or expired by IBM Spectrum Protect Snapshot according to the defined policy. The expiration does not delete an incremental backup. Two different active backups can exist at the same time:
 - Full backup, along with any associated incremental backups and differential backups.
 - Copy backup, along with any associated incremental backups and differential backups.
- When you delete an active full or copy backup, the state of the previous active full or copy backup changes from inactive to active. However, the current active incremental or differential backup is not deleted and erroneously seems to be associated with the newly active full or copy backup. Also, the incremental or differential backup (associated with the previous inactive full or copy backup that is now changed to active) remains inactive. This inactive incremental or differential backup might not display in the query output unless the **/all** parameter is specified with the **query fcm** command.
- If you delete multiple LOCAL snapshots that are stored on SAN Volume Controller, Storwize family or Space Efficient volumes, you must do so in the same order in which you created the snapshots. That is, you must delete the oldest one first, followed by the second oldest. Failure to delete them in this order can cause removal of other snapshots of the same source.
- If you mount a local VSS COPY type backup as a snap of a snap, the snap of a snap volume is also deleted along with the VSS backup.

Delete backup syntax

Use the **delete backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Delete backup positional parameters

Positional parameters immediately follow the **delete backup** command and precede the optional parameters.

The following positional parameters specify the backup to delete:

|*dbname1,...,dbnameN* *backuptype

*

Delete the active backups of all databases.

dbname

Delete a backup of the specified database. The active backup is deleted unless you specify a different backup with the **/object** parameter. When multiple active incremental backups exist, the **/object** parameter must be specified with the **delete** command.

Multiple entries are separated by commas. If separated by commas, ensure that there is no space between the comma and the component name. If any component name contains blanks, enclose the component name in double quotation marks.



Attention:

- Deleting incremental or differential backups can cause loss of recovery points.
- Deleting a full backup might cause incremental or differential backups to remain in a suspended state and are considered useless without a corresponding full backup.

The following positional parameters specify the type of delete backup to run:

FULL|COPY|INCRemental|DIFFerential

FULL

Delete full type backups.

COPY

Delete copy type backups.

INCRemental

Delete incremental type backups.

DIFFerential

Delete differential type backups.

Delete backup optional parameters

Optional parameters follow the **delete backup** command and positional parameters.

/BACKUPDESTination=LOCAL|TSM

Use this parameter to specify the destination of the backups to be deleted. The default is TSM

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values to use for a **delete backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=*nodename*

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/FROMEXCSErver=*server-name*

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was processed.

The default is the local Exchange Server.

If a DAG node is specified by using the **dagnode** parameter, IBM Spectrum Protect Snapshot for Microsoft Exchange Server uses this node name instead of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **delete** command automatically deletes the backups that are created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Exchange Server to process operations, use the **/logfile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJECT=objectname

Use the **/object** parameter to specify the name of the backup object you want to delete. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot for Exchange Server.

Use the IBM Spectrum Protect Snapshot for Exchange Server **query fcm * /all** command to view the names of all available backup objects.

The **/object** parameter is used to delete only one incremental backup at a time. When multiple active incremental backups exist, the **/object** parameter must be specified with the **delete backup** command. If it is not specified, the **delete backup** command fails.

/OLDERthan=days

Use the **/olderthan** parameter to specify how old backup files can be to be deleted. The *days* variable can range from 0 - 9999. There is no default value for the **/olderthan** parameter.

/Quiet

This parameter prevents status information from being displayed. This function does not affect the level of information that is written to the activity log.

/TSMNODE=tsmnode name

Use the *tsmnode name* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTfile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the IBM Spectrum Protect Snapshot installation directory is used.

If the *tsmoptfilename* variable includes spaces, enclose it in double quotation marks. For example:

```
/TSMOPTfile="c:\Program Files\file.opt"
```

The default is *dsm.opt* if you do not specify the **/tsmoptfile** parameter or if you specify **/tsmoptfile** but not *tsmoptfilename*.

TSMPassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS** GENERATE in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS** GENERATE is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS** PROMPT is in effect, and you do not specify a password value on the command line, you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Delete Backup example

This output example provides a sample of the text, messages, and process status that displays when you use the **delete backup** command.

In this example, the command deletes a full backup of database *rabbitvm3_sw2ie_mbdb1*. The following output is displayed:

```
Connecting to IBM Spectrum Protect Server as node 'RABBITVM3_EXCH'...
Connecting to Local DSM Agent 'RABBITVM3'...
Using backup node 'RABBITVM3_EXCH'...
Backups to be deleted:
<rabbitvm3_sw2ie_mbdb1 : VSS : full : 10/09/2014 13:30:12>
```

```
VSS Delete backup operation completed with rc = 0
Files Examined   : 1
Files Completed  : 1
Files Failed     : 0
Total Bytes      : 0
```

```
The operation completed successfully. (rc = 0)
```

Help command


Use the **tdpexcc help** command to display help for IBM Spectrum Protect Snapshot for Exchange Server commands.

This command lists one or more commands and their parameters. When you use a language other than English, you might be required to set the width of your screen display. To view the entire help description in one screen, set the screen display width to a value greater than 80 characters. For example, set the screen width to 100 characters.

Help syntax

Use the **help** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Help optional parameters

Optional parameters follow the IBM Spectrum Protect Snapshot for Exchange Server **help** command.

The following optional parameters specify the help to be displayed:

*|**command**

Identifies the specific IBM Spectrum Protect Snapshot for Exchange Server command that is to be displayed. If you specify the asterisk (*) wildcard character, help for all IBM Spectrum Protect Snapshot for Exchange Server commands are displayed.

*|**subcommand**

Help can be displayed for commands that have several subcommands, for example, the **query** command. If you do not specify a subcommand or the asterisk (*) wildcard character, help for all IBM Spectrum Protect Snapshot for Exchange Server **query** commands are displayed.

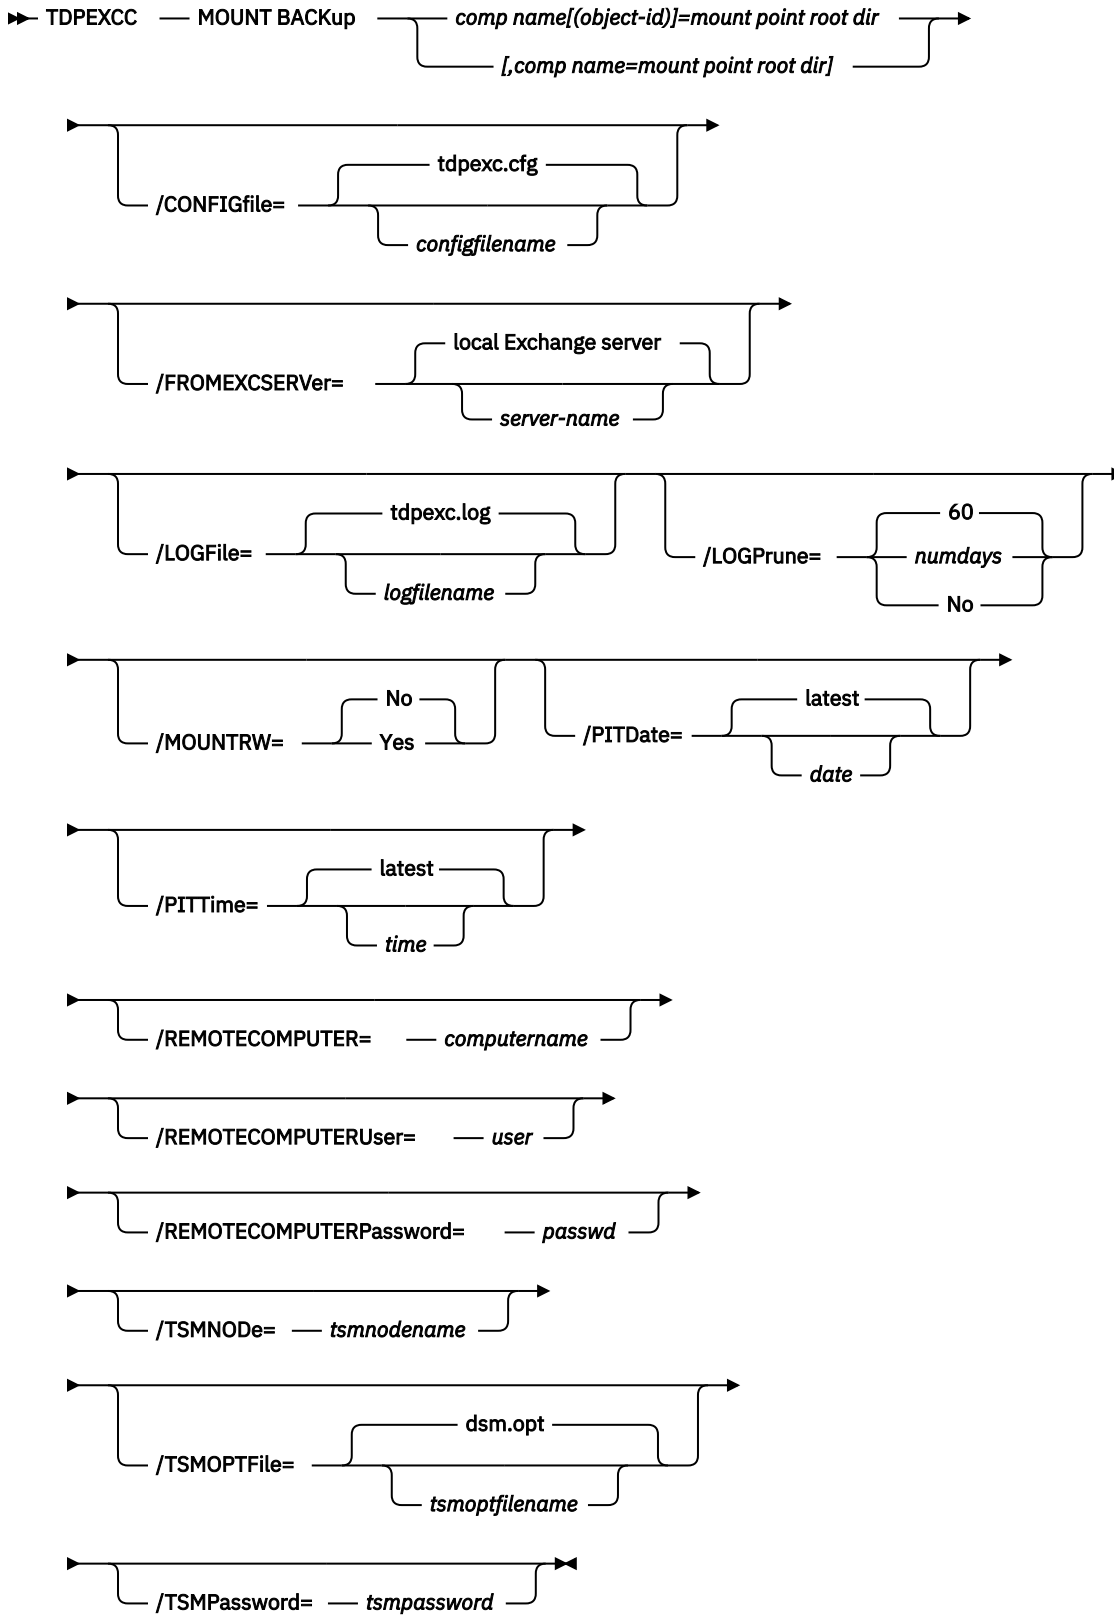
Mount backup command

To mount backups that are managed by IBM Spectrum Protect Snapshot for Exchange Server, use the **mount backup** command.

Mount backup syntax

Use the **mount backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Mount backup positional parameter

The positional parameters immediately follow the **mount backup** command and precede the optional parameters.

The following positional parameters specify the objects to mount:

component name[(object-id)]=mount point root dir[,component name=mount point root dir]

component name[(object-id)]

Specify the backup of a local Exchange database.

mount point root dir

Specify the absolute path to the directory where the snapshots are going to be displayed as mount point directories. The directory must be empty. If not empty, an error is reported.

The list must contain all non-qualified objects or all qualified objects. The list cannot contain a combination of non-qualified objects and qualified objects. Specify the list by using the following syntax:

```
mount backup object-1[(object-1-id)]= mount-point-1[,object-2[(object-2-id)]
=mount-point-2...]
```

For example:

```
tdpexc mount backup excdb(20120815064316)=f:\emptyfolder
```

Mount backup optional parameters

Optional parameters follow the **mount backup** command and positional parameters.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values to use for a **mount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpexc.cfg"
```

/FROMEXCSErver=server-name

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was processed.

The default is the local Exchange Server.

If a DAG node is specified by using the **dagnode** parameter, IBM Spectrum Protect Snapshot for Microsoft Exchange Server uses this node name instead of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **delete** command automatically deletes the backups that are created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpexc.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option No can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify no, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

Perform a read-only mount operation.

Yes

Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPMount** parameter in the configuration file.

- If **USESNAPOFASNAPMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
- If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which require IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which require IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate more target volumes on your storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

/REMOTECOMPUTER=computername

Enter the computer name or IP address of the remote system where the backup was created.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **/REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/TSMOPTFile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\dsm.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

For example:

```
TDPEXCC MOUNT BACKup EXC-DB-1=C:\MP-dir
```

```
TDPEXCC MOUNT BACKup EXC-DB-1(20120523061914)=C:\MP-dir-2 /MOUNTWritable=Yes
```

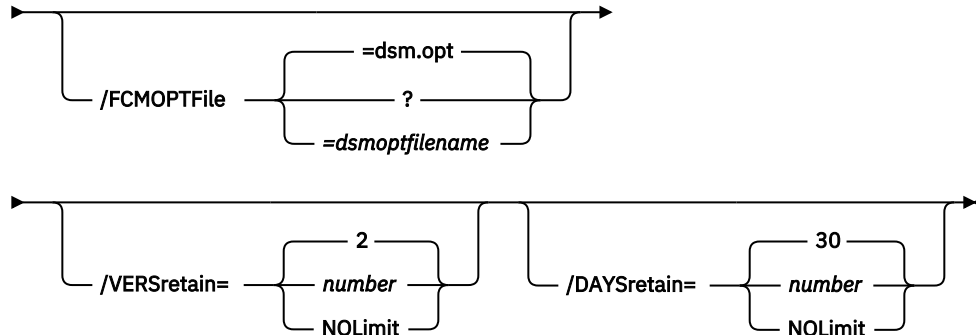
Policy commands for IBM Spectrum Protect Snapshot for Exchange

Create Policy

This command is used to create a policy.

TDPEXCC command: CREATE POLIcy

➤ TDPEXCC — CREATE POLIcy — *policy_name* →



Parameters:

- **policy_name** (required): Specifies the name of the policy that is being created. To create a policy, the policy name must be unique.
- **FCMOPTFile**: Specifies the IBM Spectrum Protect Snapshot options file to use.
- **VERSretain**: Specifies the number of snapshot versions to retain (1 - 9999). You can also specify `NOLimit` to represent an unlimited number of snapshot versions to retain.

This parameter does not apply to incremental backup versions of Exchange Server data. Incremental backups do not participate in expirations because of version limits. There is never more than one version of an incremental backup object. There is only one version of an incremental backup object because incremental backups are always uniquely named.

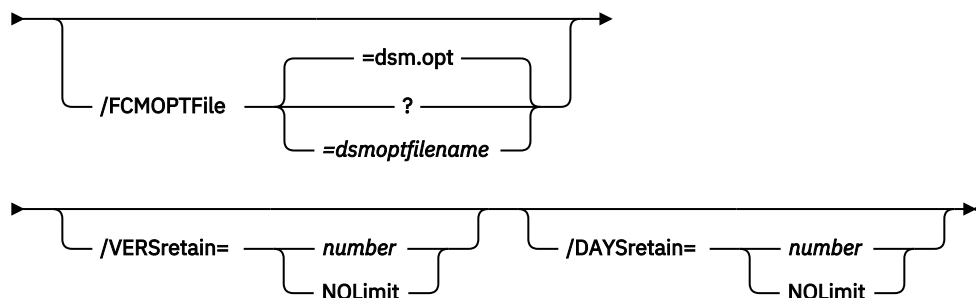
- **DAYSretain**: Specifies the number of days to retain a snapshot (0 - 9999). You can also specify `NOLimit` to represent an unlimited number of days to retain snapshot versions.

Update Policy

This command is used to update or modify an existing policy.

TDPEXCC command: UPDATE POLIcy

➤ TDPEXCC — UPDATE POLIcy — *policy_name* — `/NEWname=` — *string* →



Parameters:

- **NEWname**: Specifies the new name of the policy, if the name is being updated. The policy name must be unique.
- **policy_name** (required): Specifies the name of the policy that is being updated.
- **FCMOPTFile**: Specifies the IBM Spectrum Protect Snapshot options file to use.
- **VERSretain**: Specifies the number of snapshot versions to retain (1 - 9999). You can also specify NOLimit to represent an unlimited number of snapshot versions to retain.

This parameter does not apply to incremental backup versions of Exchange Server data. Incremental backups do not participate in expirations because of version limits. There is never more than one version of an incremental backup object. There is only one version of an incremental backup object because incremental backups are always uniquely named.


- **DAYSretain**: Specifies the number of days to retain a snapshot (0 - 9999). You can also specify NOLimit to represent an unlimited number of days to retain snapshot versions.

Copy Policy

This command is used to copy an existing policy to a new policy.

TDPEXCC command: COPY POLicy

▶ TDPEXCC — COPY POLicy — *existing_policy_name* — *new_policy_name* →



Parameters:

- **existing_policy_name** (required): Specifies the name of the policy that is being copied.
- **FCMOPTFile**: Specifies the IBM Spectrum Protect Snapshot options file to use.
- **new_policy_name** (required): Specifies the name of the new policy. The policy name must be unique.

Query Policy

This command is used to list the attributes of a policy.

TDPEXCC command: Query POLicy

▶ TDPEXCC — Query POLicy — * →

Parameters: * (required) Specifies all policies are to be queried. The result of the query is displayed as follows:

```


Connecting to Exchange Server, please wait...
Policy      Number of snapshots to keep  Days to keep a snapshot
-----
FCMPOL      3                            60
STANDARD    2                            30

```

Delete Policy

This command is used to delete a policy.

TDPEXCC command: DElete POLicy



Parameters:

- **policy_name** (required): Specifies the name of the policy that is being deleted.
- **FCMOPTFile**: Specifies the IBM Spectrum Protect Snapshot options file to use.

Exchange policy examples

These output examples provide a sample of the text, messages, and process status that displays when you use the **create policy** and **delete policy** commands.

In this example, the `tdpexcc create policy fcmexchol1` command creates the *FCMEXCHPOL1* policy. The following output is displayed:

```
C:\PROGRA~1\Tivoli\tsm\TDPEXchange>tdpexcc create policy fcmexchol1

IBM Spectrum Protect for Mail:
Data Protection for Microsoft Exchange Server
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 1998, 2016. All rights reserved.

CREATE policy was successful.

The operation completed successfully. (rc = 0)
```

In this example, the `tdpexcc delete policy fcmexchol1` command deletes the *FCMEXCHPOL1* policy. The following output is displayed:

```
C:\PROGRA~1\Tivoli\tsm\TDPEXchange>tdpexcc delete policy fcmexchol1

IBM Spectrum Protect for Mail:
Data Protection for Microsoft Exchange Server
Version 7, Release 1, Level 0
(C) Copyright IBM Corporation 1998, 2016. All rights reserved.

DELETE policy was successful.

The operation completed successfully. (rc = 0)
```

Query Exchange command

Use the **query exchange** command to query the local Exchange Server for general information.

The **query exchange** command returns the following information:

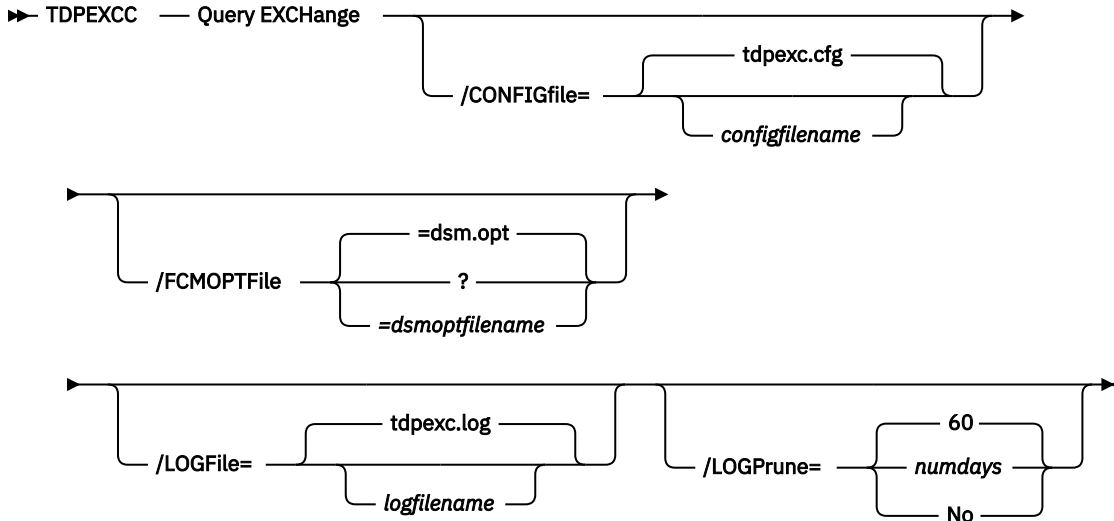
- Exchange Server name and version
- Domain name
- Names of all databases
- Status (online, offline) of all databases
- Circular logging status (enabled, disabled) of all databases
- VSS information:
 - Writer Name
 - Local DSMAgent Node
 - Remote DSMAgent Node

- Writer Status (online, offline)
- Number of selectable components

Query Exchange syntax

Use the **query exchange** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Query Exchange optional parameters

Optional parameters follow the **query exchange** command.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values to use for a **query exchange** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/FCMOPTFile=dsmoptfilename

The **/FCMOPTFile** parameter specifies the IBM Spectrum Protect Snapshot options file to use.

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/FCMOPTFile**, the default value is `dsm.opt`.
- If you specify **/FCMOPTFile** but not *dsmoptfilename*, the default is also `dsm.opt`.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server. The *logfilename* variable identifies the name of the activity log file. If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for

Exchange Server installation directory. If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *tdpexc.log*. The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

Query Exchange example

This output example provides a sample of the text, messages, and process status that displays when you use the **query exchange** command.

The **tdpexcc query exchange** command queries the Exchange server. An example of the output in an Exchange Server Database Availability Group environment is displayed:

Querying Exchange Server to gather component information, please wait...

Microsoft Exchange Server Information

```
-----  
Server Name:          AVATAR  
Domain Name:         avatar.local  
Exchange Server Version: 14.1.270.1
```

Databases and Status

```
-----  
  
avatarDB3_D_local_bas  
Circular Logging - Disabled  
DAG Status - None  
Recovery - False  
avatarDB3_D_local_bas           Offline  
  
avatarDB4_D_local_bas  
Circular Logging - Disabled  
DAG Status - None  
Recovery - False  
avatarDB4_D_local_bas           Online  
  
avatarDB5_G_storwize_bas  
Circular Logging - Disabled  
DAG Status - None  
Recovery - False  
avatarDB5_G_storwize_bas       Online  
  
avatar_F_H  
Circular Logging - Disabled  
DAG Status - None  
Recovery - False  
avatar_F_H                       Online  
  
Mailbox Database 0003208508  
Circular Logging - Disabled  
DAG Status - None  
Recovery - False  
Mailbox Database 0003208508     Online
```

Volume Shadow Copy Service (VSS) Information

```
-----  
Writer Name           : Microsoft Exchange Writer  
Local DSMAgent Node   : AVATAR  
Remote DSMAgent Node  :  
Writer Status         : Online  
Selectable Components : 4
```

Query FCM command

Use the **query fcm** command to display IBM Spectrum Protect Snapshot information.

This command displays the following information:

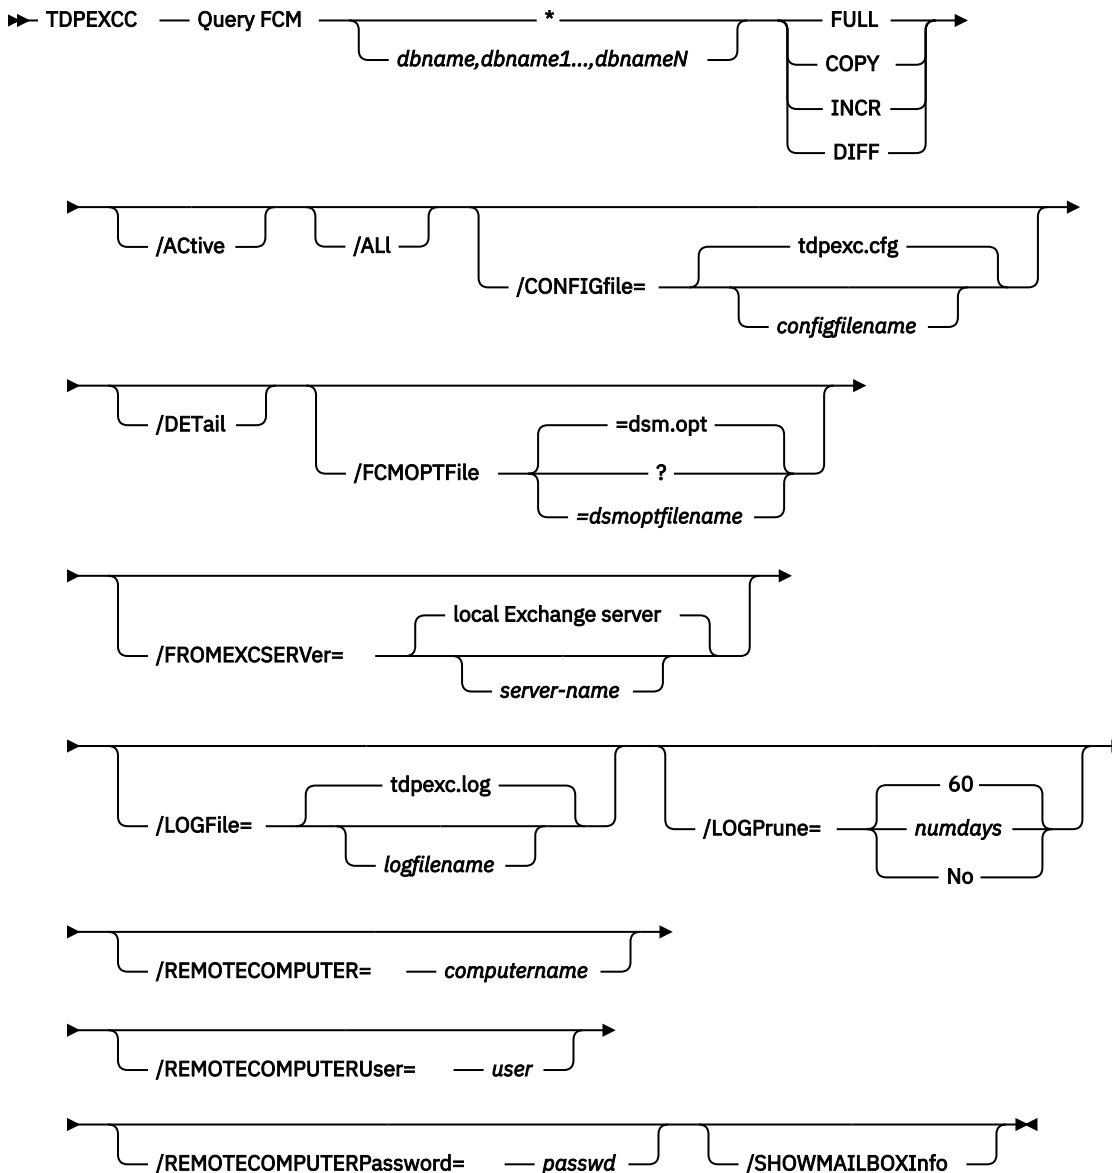
- Compression mode
- Active policy set
- Default management class

This command can also display a list of backups that match the databases that are entered.

Query FCM syntax

Use the **query FCM** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Query FCM positional parameters

Positional parameters immediately follow the **query FCM** command and precede the optional parameters.

The following positional parameters specify the object to query. If none of these positional parameters are specified, only the IBM Spectrum Protect Snapshot API and IBM Spectrum Protect Snapshot information is displayed:

* | *dbname*

dbname1, ..., dbnameN

Query all backup objects for the specified database. Multiple entries are separated by commas. where *dbname* can be a database name.

The following positional parameters specify the type of backup to query. If this parameter is not specified, all backup types are displayed:

FULL

Query only full backup types.

COPY

Query only copy backup types.

INCR

Query only incremental backup types.

DIFF

Query only differential backup types.

Query FCM optional parameters

Optional parameters follow the **query FCM** command and positional parameters.

/Active

Use the **/Active** parameter to display active backup objects only. This parameter is the default.

/All

Use the **/All** parameter to display both active and inactive backup objects. If the **/All** parameter is not specified, only active backup objects are displayed.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values for the IBM Spectrum Protect Snapshot for Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/Detail

Use the **/Detail** parameter to display detailed output from the query command.

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/FCMOPTFile**, the default value is `dsm.opt`.
- If you specify **/FCMOPTFile**, but not *dsmoptfilename*, the default is also `dsm.opt`.

/FCMOPTFile=dsmoptfilename

The **/FCMOPTFile** parameter specifies the IBM Spectrum Protect Snapshot options file to use.

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/FCMOPTFile**, the default value is `dsm.opt`.
- If you specify **/FCMOPTFile**, but not *dsmoptfilename*, the default is also `dsm.opt`.

/FROMEXCServer=server-name

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was processed.

The default is the local Exchange Server.

If a DAG node is specified by using the **dagnode** parameter, IBM Spectrum Protect Snapshot for Microsoft Exchange Server uses this node name instead of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server node to back up databases in an Exchange Server Database Availability

Group. Therefore, the **delete** command automatically deletes the backups that are created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *tdpexc.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to query the data that is backed up.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/SHOWMAILBOXInfo

Use the **/SHOWMAILBOXInfo** parameter to display mailbox history information in backup databases.

Query FCM example

The following command shows detailed information about current backups: **query fcm * /detail**

```
IBM FlashCopy Manager for Mail:
FlashCopy Manager for Microsoft Exchange Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

Querying FlashCopy Manager server for a list of database backups, please wait...

Connecting to FCM Server as node 'TIVVM483_EXC'...
Connecting to Local DSM Agent 'TIVVM483'...
Using backup node 'DAG2'...

DAG : DAG2

Backup Object Information
-----

Exchange Server Name ..... TIVVM483
Database Availability Group ..... DAG2
Backup Database Name ..... RATTEST_DAGDB
Backup Method ..... VSS
Backup Location ..... Loc
Backup Object Type ..... full
Mount Points Root Directory .....
Backup Object State ..... Active
Backup Creation Date / Time ..... 08/22/2014 22:23:00
Backup Supports Instant Restore ..... No
Backup Object Size / Name ..... 172.07MB / 20120822222300
Backup Object Size / Name ..... 36.01MB / Logs
Backup Object Size / Name ..... 136.06MB / File

The operation completed successfully. (rc = 0)
```

Query Managedcapacity command

Use the **Query Managedcapacity** command to assist with storage planning by determining the amount of managed capacity in use.

Purpose

The **Query Managedcapacity** command displays capacity-related information about the volumes that are represented in local inventory that is managed by IBM Spectrum Protect Snapshot. This command is valid for all Windows operating systems that are supported by IBM Spectrum Protect Snapshot.

TDPEXCC command

► TDPEXCC — Query MANAGEDCAPacity 

Parameters

/Detailed

Results in a detailed listing of snapped volumes. If this option is not specified, then only the total capacity is displayed.

In this example, the **tdpexcc query managedcapacity** command displays the total amount of managed capacity in use in the local inventory. The following output is displayed:

```
Total Managed Capacity : 47.99 GB (51,533,307,904 bytes)
```

In this example, the **tdpexcc query managedcapacity /detailed** command displays a detailed listing of total amount of managed capacity and the snapped volumes in use. The following output is displayed:

```
IBM FlashCopy Manager for Mail:  
FlashCopy Manager for Microsoft Exchange Server  
Version 7, Release 1, Level 3.0  
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.
```

```
Total Managed Capacity : 31.99 GB (34,353,438,720 bytes)
```

```
Volume      : M:  
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

```
Volume      : F:  
Managed Capacity : 16.00 GB (17,176,719,360 bytes)  
Total Managed Capacity : 1,019.72 MB (1,069,253,632 bytes)
```

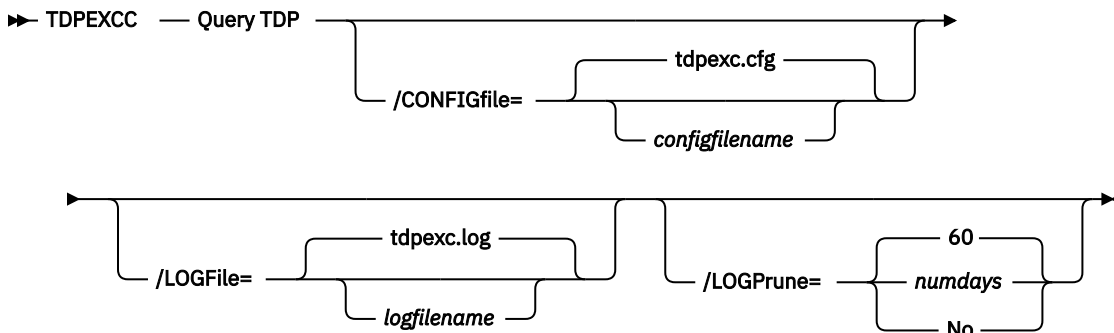
Query TDP command

Use the **query tdp** command to query a list of the current values that are set in the configuration file for IBM Spectrum Protect Snapshot for Exchange Server.

Query TDP syntax

Use the **query TDP** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Query TDP optional parameters

Optional parameters follow the **query TDP** command.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values to use for a **query tdp** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *tdpexc.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option *No* can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

Query TDP example

This output example provides a sample of the text, messages, and process status that displays when you use the **query TDP** command.

An example of the output in a VSS configuration is displayed.

```
IBM FlashCopy Manager for Mail:
FlashCopy Manager for Microsoft Exchange Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1998, 2015. All rights reserved.

FlashCopy Manager for Exchange Preferences
-----

BACKUPDESTination..... LOCAL
DATEformat ..... 1
IMPORTVSSSNAPSHOTSONLYWhenneeded ... No
LANGuage ..... ENU
LOCALDSMAgentnode..... CENTORI
LOGFile ..... tdpexc.log
LOGPrune ..... 60
MOUNTWait ..... Yes
NUMBERformat ..... 1
REMOTEDSMAgentnode.....
TEMPDBRestorepath.....
TEMPLOGRestorepath.....
TIMEformat ..... 1
```

Restore command

Use the **restore** command to restore a database backup from local shadow volumes that are managed by IBM Spectrum Protect Snapshot to an Exchange Server.

To complete an IBM Spectrum Protect Snapshot for Exchange Server restore, you must have local registry rights for all versions of the Exchange Server.

When you use the **restore** command, remember the following guidelines:

- When you restore inactive backups or active incremental backups, use the **/object** parameter to specify the name of the backup object to restore. This object name uniquely identifies the backup instance that is managed by IBM Spectrum Protect Snapshot storage. You can enter a **tdpexcc query fcm * /all** command to obtain a list of all the active and inactive backup objects.

If the `tdpexcc restore dbname incr` command is entered (without the **/object** parameter) to restore multiple active incremental backups, all multiple active incremental backups are restored sequentially. The **/object** parameter is used to restore only one incremental backup at a time.

- Use the **/eraseexistinglogs** parameter to direct the program to erase the existing transaction log files for the database before it restores the database. If you do not specify this option, existing transaction logs are not erased, and might be reapplied when the Exchange databases are mounted. This parameter is only valid when you restore a FULL or COPY VSS backup of Exchange Server databases.
- Specify **/mountdatabases=yes** if you are restoring the last backup and you want the database or databases to be automatically mounted after the recovery is completed. Only transaction logs that are contained in the backup is applied to the mailbox database when you run a recovery database restore. You must specify **/recover=applyrestoredlogs** when you restore a mailbox database to a recovery database. Otherwise, the restore operation might fail.

The graphical user interface provides an easy-to-use, flexible interface to help you run a restore operation. The interface presents information in a way that allows multiple selection and, in some cases, automatic operation.

With Microsoft Exchange Server, you cannot specify the asterisk (*) wildcard character in database names. Databases that contain the asterisk (*) wildcard character in their name are not backed up.

IBM Spectrum Protect Snapshot for Exchange Server supports the following types of restore:

Full

Restore a full type backup.

Copy

Restore a copy type backup.

Incremental

Restore an incremental type backup.

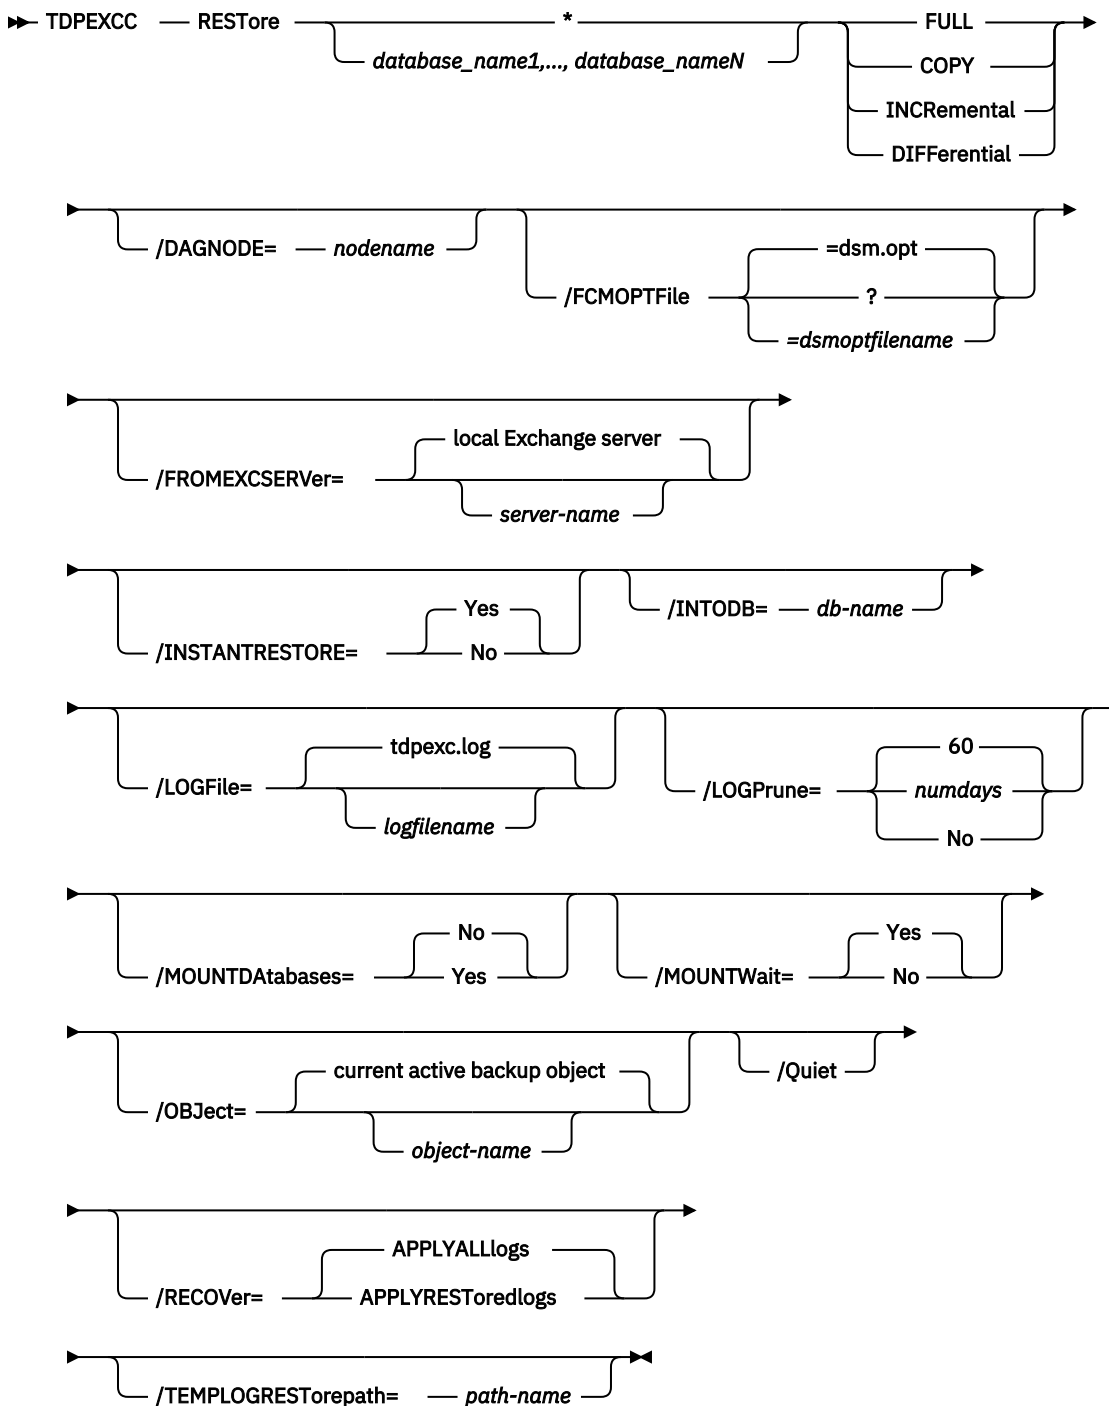
Differential

Restore a differential type backup.

Restore syntax

Use the **restore** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Restore positional parameters

Positional parameters immediately follow the **restore** command and precede the optional parameters.

The following positional parameters specify the object to restore:

*** | *database_name1, ..., database_nameN***

Restore all components sequentially.

The following positional parameters specify the type of restore to run:

FULL | COPY | INCRemental | DIFFerential

FULL

Restore a full backup.

COPY

Restore a copy backup.

INCRemental

Restore an incremental backup.

DIFFerential

Restore a differential backup.

Restore optional parameters

Optional parameters follow the **restore** command and positional parameters.

/CONFIGfile=configfilename

Use the ***/CONFIGfile*** parameter to specify the name of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file that contains the values for the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Microsoft Exchange Server installation directory is used. If the ***/CONFIGfile*** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire ***/CONFIGfile*** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/DAGNode=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/EXCLUDEDUMPster=dsmoptfilename

/FCMOPTfile=dsmoptfilename

The ***/FCMOPTfile*** parameter specifies the IBM Spectrum Protect Snapshot options file to use.

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify ***/FCMOPTfile***, the default value is `dsm.opt`.
- If you specify ***/FCMOPTfile*** but not *dsmoptfilename*, the default is also `dsm.opt`.

/FROMEXCServer=server-name

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was processed.

The default is the local Exchange Server.

If a DAG node is specified by using the **dagnode** parameter, IBM Spectrum Protect Snapshot for Microsoft Exchange Server uses this node name instead of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **delete** command automatically deletes the backups that are created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/INSTANTRestore=Yes|No

Use the **/INSTANTRestore** parameter to specify whether to use volume level snapshot or file level copy to restore a VSS backup that is on local shadow volumes. The default value is Yes. An IBM Systems Storage SAN Volume Controller, DS8000, the XIV system, and IBM Storwize V7000 storage system is required to complete VSS instant restore operations.

You can specify:

Yes

Use volume level snapshot restore for a VSS backup that is on local shadow volumes if the backup exists on volumes that support it. This option is the default.

No

Use file level copy to restore the files from a VSS backup that is on local shadow volumes. Bypassing volume-level copy means that Exchange database files, log files, and the checkpoint file are overwritten on the source volumes.

When a VSS instant restore is completed on DS8000 and Storwize family, make sure that any previous background copies (that involve the volumes that are being restored) are completed before you initiate the VSS instant restore operation. The **/instantrestore** parameter is ignored and VSS instant restore capabilities are automatically disabled when it runs any type of VSS restore into operation. VSS instant restore of differential and incremental backups is not supported.

/INTODB=db-name

Use the **/INTODB** parameter to specify the name of the database into which the VSS backup is restored. The database name must be specified with the *db-name* variable. For example, if RDB is the name of the database into which the VSS backup is restored, the input on the command line is as follows:

```
TDPEXCC RESTore Maildb1 FULL /INTODB=RDB
```

However, when you restore a database that is relocated (system file path, log file path, or database file path), you must specify the same database name as the one you are restoring. For example, if *Maildb5* is the name of the relocated database that is being restored, the command-line entry is as follows:

```
TDPEXCC RESTore Maildb5 FULL /INTODB=Maildb5
```

- There is no default value.
- To restore into a Recovery Database (RDB) or alternate database, an RDB or alternate database must exist before you attempt the restore operation.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Microsoft Exchange Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Microsoft Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If you do not specify the **/LOGFile** parameter, log records are written to the default log file, *tdpexc.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Microsoft Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option *No* can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTDatabases=No|Yes

Use the **/mountdatabases** parameter to specify whether to mount the databases after the restore operation is completed. You must specify one of the following values:

Yes

Mount the databases after the restore operation is completed.

No

Do not mount the databases after the restore operation is completed. This option is the default.

/MOUNTWait=Yes | No

Use the **/mountwait** parameter to specify whether IBM Spectrum Protect Snapshot for Microsoft Exchange Server waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify these options:

Yes

IBM Spectrum Protect Snapshot for Microsoft Exchange Server waits until all initial volumes of any required removable media are made available to the IBM Spectrum Protect server before it completes the command. This option is the default.

No

IBM Spectrum Protect Snapshot for Microsoft Exchange Server ends the command (if removable media are required). An error message is displayed.

/OBJECT=object-name

Use the **/object** parameter to specify the name of the backup object you want to restore. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot for Microsoft Exchange Server.

Use the IBM Spectrum Protect Snapshot for Microsoft Exchange Server query `fcM /all` command to view the names of active and inactive backup objects.

If the `tdpexcc restore dbname incr` command is entered (without the **/object** parameter) to restore multiple active incremental backups, all multiple active incremental backups are restored sequentially. The **/object** parameter is used to restore only one incremental backup at a time.

/Quiet

This parameter prevents status information from being displayed. This function does not affect the level of information that is written to the activity log.

/RECOVER=APPLYRESToredlogs|APPLYALLlogs

Use this parameter to specify whether you want to run recovery after you restore an object. If the database is not mountable, you can either restore the last backup again and specify the **/RECOVER=value** option or you can use the Microsoft **ESEUTIL /cc** command to run recovery manually.

You must specify one of the following values when you use this parameter:

APPLYALLlogs

Specify `/recover=applyalllogs` to replay the restored-transaction log entries and the current active-transaction log entries. Any transaction logs entries that display in the current active-transaction log are replayed. This option is the default.

APPLYRESToredlogs

Specify `/recover=applyrestoredlogs` to replay only the restored-transaction log entries. The current active-transaction log entries are not replayed. When you choose this option for a restore, your next backup must be a full or copy backup.

When you restore multiple backup objects, the **/RECOVER** option must be used on the restore of the last object.

/TEMPLOGRESTorepath=path-name

Use the **/TEMPLOGRESTorepath** parameter to specify the default temporary path to use when logs and patch files are restored. For best performance, this path must be on a different physical device than the current active-transaction logger.

If you do not specify the **/TEMPLOGRESTorepath** parameter, the default value is the value that is specified by the **/TEMPLOGRESTorepath** option in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file. The default IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file is `tdpexc.cfg`.

If you do not specify the **/TEMPLOGRESTorepath** parameter, and the **/TEMPLOGRESTorepath** value does not exist in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file, the TEMP environment variable value is used.

When you do a full or copy restore operation, all log files in the path that is specified by the **/TEMPLOGRESTorepath** parameter are erased. In addition, the value of **/TEMPLOGRESTorepath** must not be the same value as the current location for the database. If the value is the same, the database can become corrupted.

Do not specify double-byte characters (DBCS) within the temporary log path.

Restore example

This output example provides a sample of the text, messages, and process status that displays when you use the **restore** command.

In this example, the command completes an instant restore of the local backup for mailbox database *rabbitvm3_sw2ie_mbdb1*. The following output is displayed:

```
Connecting to TSM Server as node 'RABBITVM3_EXCH'...
Connecting to Local DSM Agent 'RABBITVM3'...
Using backup node 'RABBITVM3_EXCH'...
Starting Microsoft Exchange restore...

Beginning VSS restore of 'rabbitvm3_sw2ie_mbdb1'. This operation could take a while,
please wait...
Restoring 'rabbitvm3_sw2ie_mbdb1' via volume-level copy from snapshot(s). This process may
take some time. Please wait.
VSS Restore operation completed with rc = 0
  Files Examined      : 0
  Files Completed     : 0
  Files Failed        : 0
  Total Bytes         : 0
  Total LanFree Bytes : 0

Running recovery. This operation might take some time, depending on the number
of transaction logs being replayed.

The operation completed successfully. (rc = 0)
```

Restorefiles command

Use the **restorefiles** command to restore flat files from a backup into a specified directory.

The following information provides details about this using the **restorefiles** command:

- The **restorefiles** command is only available on the command-line interface.
- This command does not require an Exchange Server to be installed on, or accessible from the system where **restorefiles** is run.
- Files can be restored to an alternative system or to an alternative directory on the same system as the Exchange Server.
- The **restorefiles** operation fails if a previously restored file exists, except for VSS backup files.
- The command continues until it succeeds, or until the destination volume does not contain enough space for the operation.
- When you restore files from an inactive backup or an active incremental backup, use the **/object** parameter to specify the name of the backup object. The object name uniquely identifies the backup instance in IBM Spectrum Protect server storage. A list of backup object names is obtained by issuing the **query tsm * /all** command.

A VSS **restorefiles** operation overwrites files that exist and have the same name. If a log file from an incremental backup has the same name as the log file from the full backup operation, you can run two consecutive **restorefiles** operations to the same directory:

- Run the following command to restore a full backup:

```
tdpexcc restorefiles DB1 FULL /into=d:\temprestore
```

- Run the following command to restore the log files during the incremental restore:

```
tdpexcc restorefiles DB1 INCR /into=d:\temprestore
```

The **/into** parameter can be a directory on any mounted local disk. It is not possible to use a mapped network drive as a restore destination.

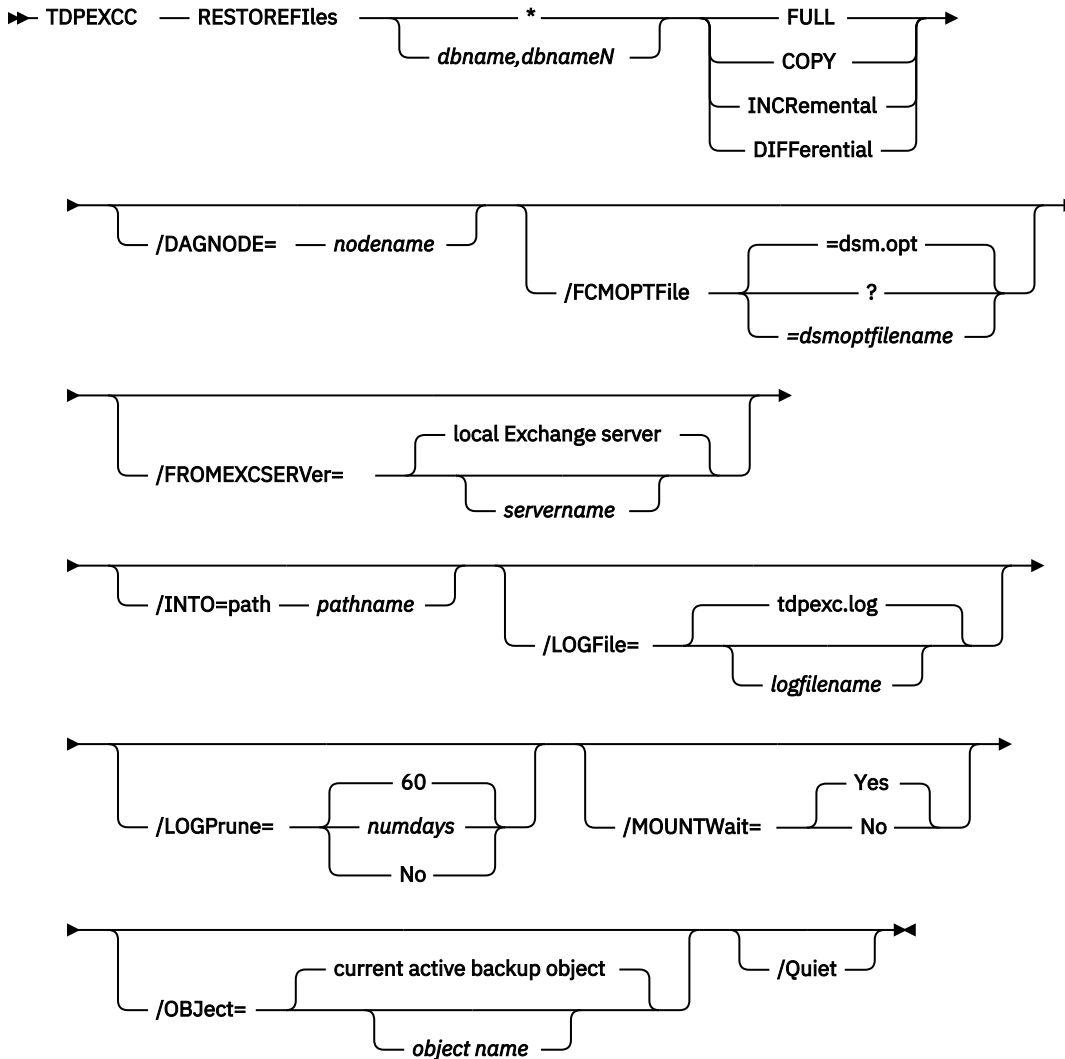
Before you issue the **restorefiles** command, make sure that you have sufficient disk space to hold all of the flat files. For example, if your database and logs are 50 GB in size, you need 50 GB available in the

destination directory that is specified by the **/into** parameter. For VSS backups, do not issue a **restorefiles** command to the existing location of the production or active database. Those files are overwritten.

Restorefiles syntax

Use the **restorefiles** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Restorefiles positional parameters

Positional parameters immediately follow the **restorefiles** command and precede the optional parameters.

The following positional parameters specify the object to restore:

* *dbname*

*

Sequentially restore all flat files for the database.

dbname

Restore the specified database files. Multiple entries are separated by commas.

The following positional parameters specify the type of backup from which the files are restored:

FULL | COPY | INCRemental | DIFFerential *dbname*

FULL

Restore the files from a full backup.

COPY

Restore the files from a copied backup.

INCRemental

Restore the files from an incremental backup.

DIFFerential

Restore the files from a differential backup.

Restorefiles optional parameters

The optional parameters for the **restorefiles** command and positional parameters are listed.

/CONFIGfile=*configfilename*

Use the **/CONFIGfile** parameter to specify the name of the IBM Spectrum Protect Snapshot for Exchange Server configuration file that contains the values for the IBM Spectrum Protect Snapshot for Exchange Server configuration options.

The *configfilename* variable can include a full path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=*nodename*

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/FROMEXCSErver=*server-name*

Use the **/fromexcserver** parameter to specify the name of the Exchange Server where the original backup was processed.

The default is the local Exchange Server.

If a DAG node is specified by using the **dagnode** parameter, IBM Spectrum Protect Snapshot for Microsoft Exchange Server uses this node name instead of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server node to back up databases in an Exchange Server Database Availability Group. Therefore, the **delete** command automatically deletes the backups that are created by the other DAG members, without having to specify the **/fromexcserver** parameter.

/INTO=*path*

Use the **/INTO** parameter to specify the root directory where files are to be restored. The **restorefiles** operation creates a subdirectory under the root directory that contains the name of the database. Restored files are placed in that subdirectory. If the **/INTO** parameter is not specified, the files are restored into the directory where the **restorefiles** command is issued. The **/into** parameter can be a directory on any mounted local disk. It is not possible to use a mapped network drive as a restore destination.

/LOGFile=*logfile*

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTWait=Yes | No

Use the **/mountwait** parameter to specify whether IBM Spectrum Protect Snapshot for Microsoft Exchange Server waits for removable media to mount (such as tapes or CDs) or to stop the current operation. This situation occurs when the is configured to store backup data on removable media and waits for a required storage volume to be mounted.

You can specify these options:

Yes

IBM Spectrum Protect Snapshot for Microsoft Exchange Server waits until all initial volumes of any required removable media are made available to the IBM Spectrum Protect server before it completes the command. This option is the default.

No

IBM Spectrum Protect Snapshot for Microsoft Exchange Server ends the command (if removable media are required). An error message is displayed.

/OBJECT=object

Use the **/OBJECT** parameter to specify the name of the backup object files that you want to restore. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot for Exchange Server.

Use the IBM Spectrum Protect Snapshot for Exchange Server **query tsm * /all** command to view the names of the backup objects.

/Quiet

This parameter prevents status information from being displayed. The level of information that is written to the activity log is not affected.

/FCMPTFile=dsm.opt filename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect Snapshot for Exchange Server options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot for Exchange Server is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire *dsm.opt filename* parameter entry in double quotation marks. For example:

```
/fcmptfile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

Restorem mailbox command

To restore mailbox-level data or mailbox-item-level data from IBM Spectrum Protect Snapshot for Exchange Server backups, use the **restorem mailbox** command.

The **restorem mailbox** command applies to any IBM Spectrum Protect Snapshot for Microsoft Exchange Server VSS backup:

- VSS backups that are stored on IBM Spectrum Protect server
- VSS backups that are stored on local shadow volumes

When you use the **restorem mailbox** command, follow these guidelines:

- Ensure that you have the required role-based access control (RBAC) permissions to complete individual mailbox restore operations.
- You can restore multiple mailboxes in a single mailbox restore operation.
- You can use the **restorem mailbox** command to restore data to a mailbox on the Exchange Server or to an Exchange Server .pst file.

When you restore to a Unicode .pst file, except for the **Folder Name** and **All Content** filters, the filters are ignored. The amount of time that is needed to complete the restore process depends on the size of the mailbox databases, the network speed, and the number of mailboxes to process.

For non-Unicode .pst files for Exchange Server 2013, you can limit the range of the mailbox data to restore by using the **/mailboxfilter** parameter to specify filters that are based on the following mailbox message elements:

- Sender name
- Folder name
- Message body
- Subject line
- Attachment name
- Range of the message delivery date and time

- You can use the **restoremailbox** command on the primary Exchange Server or on an alternate Exchange Server that is in the same domain.
- You can use the **restoremailbox** command to restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
 - To restore an Exchange 2013 public folder mailbox, you must have the Public Folders management role.
 - You can restore a public folder mailbox only to an existing public folder mailbox that is on the Exchange server.
 - You can restore a public folder only to an existing public folder. The public folder on the Exchange server must have the same folder path as the public folder to be restored. If the public folder is deleted from the public folder mailbox on the Exchange server, you must re-create the public folder with the same folder path as the public folder to be restored, before you start the restore operation.
 - As a best practice, restore public folder mailboxes separately from user mailboxes. Select only one public folder mailbox to restore at a time if you want to restore a specific public folder in the mailbox, or if you want to restore to a different public folder mailbox than the original mailbox.

If you restore multiple mailboxes in a single restore operation, and at least one of the mailboxes is a public folder mailbox, the mailboxes are restored only to their original mailbox locations. You cannot specify a filter or an alternate mailbox destination.
 - You can restore to a different public folder mailbox than the original mailbox if, for example, the public folder was relocated since the time of the backup. Before you complete the public folder restore operation, ensure that the public folder exists with the same folder path in the alternate mailbox location.
- In Exchange Server 2013 or later versions, you can use the **restoremailbox** command to restore an archive mailbox or only a part of the mailbox, for example, a specific folder. You can restore archive mailbox messages to an existing mailbox on the Exchange server, to an archive mailbox, or to an Exchange Server .pst file.
- You can use the **restoremailbox** command with the following parameter and options:
 - Set the **/KEEPRDB** parameter option to Yes to retain a recovery database after one or more mailboxes are restored. Set the parameter value to No to automatically remove the recovery database after mailbox restore processing.

Regardless of the option that you set, IBM Spectrum Protect Snapshot for Microsoft Exchange Server retains the recovery database if the mailbox restore operation fails after the recovery database is successfully restored.

If you restore multiple mailboxes, and you want to retain the recovery database after the restore operation is complete, ensure that all mailboxes are in the same recovery database.
 - Set the **/USEEXISTINGRDB** parameter option to Yes to restore a mailbox from an existing recovery database. Set the parameter value to No to automatically remove the existing recovery database during mailbox restore processing.
 - Set the **/mailboxoriglocation** parameter to specify the server and the database where the mailbox is at the time of backup. You set this option when the mailbox history is disabled and when the mailbox that you are restoring is either moved or deleted since the time of the backup.
 - If a mailbox is deleted or re-created since the time of the backup, you must use a temporary mailbox with enough capacity to contain all of the mailbox items that you are restoring. The mailbox of the user who is logged in is used as temporary mailbox by default.

You can set the **/tempmailboxalias** optional parameter by selecting **Properties** from the Actions pane. In the **Data Protection Properties** window, select the **General** page, where you can specify the temporary log restore path, the temporary database restore path and the alias of the temporary mailbox.
- You can use the **restoremailbox** command to recover and restore different types of mail items in the Recoverable Items folder.

- The mail items that you can restore depends on whether the mailbox is enabled for mailbox restore operations.
- You cannot restore the Recoverable Items folder and subfolder hierarchy to a mailbox restore destination. You can restore only the contents of the email folders.
- You cannot add a subfolder to the Recoverable Items folder in a mailbox.

Related conceptsExchange mailbox restore operations

By using IBM Spectrum Protect Snapshot, you can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode or non-Unicode .pst file.

Security requirements for backup and restore operations

For IBM Spectrum Protect Snapshot security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

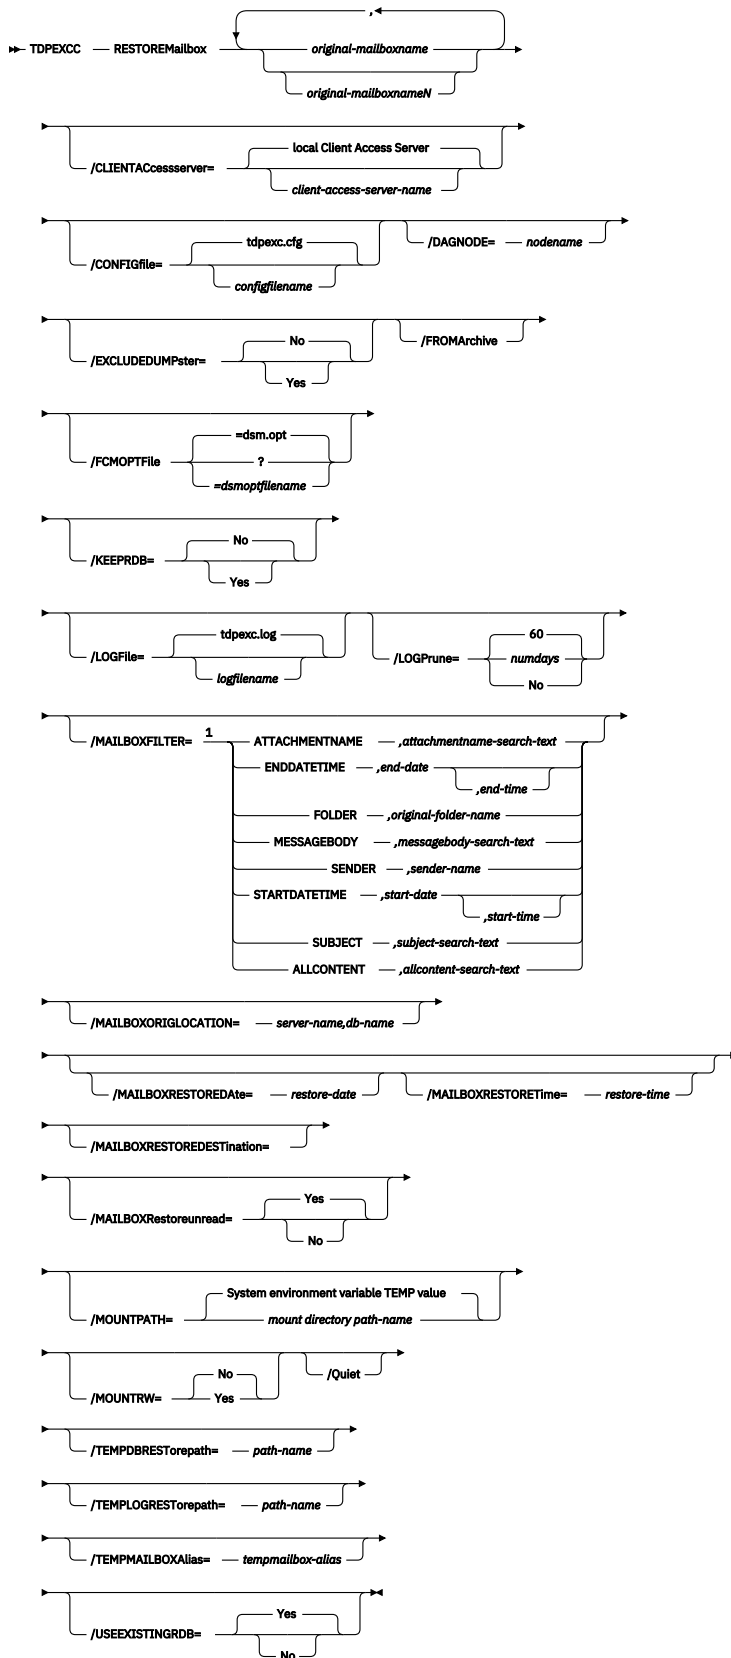
Related tasksRestoring mailbox data

IBM Spectrum Protect Snapshot backs up mailbox data at the database level, and also restores individual mailbox items from the database backup.

Restoremailbox syntax

Use the **restoremailbox** command syntax diagram as a reference to view available options and truncation requirements.

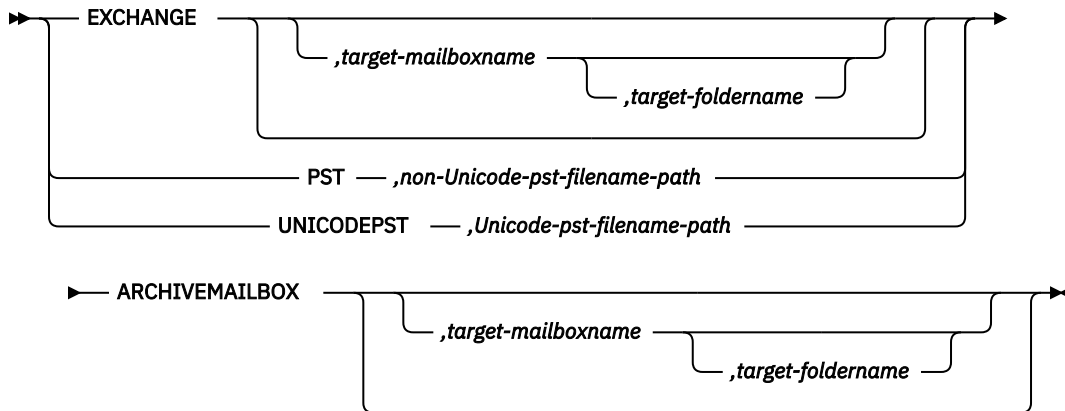
TDPEXCC command



Notes:

- ¹ You can specify the **/MAILBOXFILTER** parameter multiple times; however, you must specify each **/MAILBOXFILTER** subparameter only once.

/MAILBOXRESTOREDESTINATION options



Restoremmailbox positional parameters

Positional parameters immediately follow the `restoremmailbox` command and precede the optional parameters.

original-mailboxname

Use this parameter to specify the name of the mailbox to restore from. The mailbox name can be either the mailbox-alias, the mailbox-display name, or the mailbox globally unique identifier (GUID). The *original-mailboxname* parameter is required.

To specify more than one name, separate them by commas.

If any mailbox name contains blanks, enclose the entire mailbox name in double quotation marks.

Restoremmailbox optional parameters

Optional parameters are supplied following the `restoremmailbox` command and positional parameters.

/CLIENTAccessserver=configfilename

Use the `/CLIENTAccessserver` parameter to specify the name of the Client Access Server (CAS) that you want to use. This parameter is available only if you use Microsoft Exchange 2013 or later versions.

By default, IBM Spectrum Protect Snapshot uses the local server as the CAS if the CAS role is installed on the local server. If the CAS role is not installed on the local server, IBM Spectrum Protect Snapshot uses the mailbox database that the user is logged in to.

To determine the name of the CAS in use, run this Exchange Management Shell command:

```
Get-MailboxDatabase -Identity <logon user mailbox database> |  
select RpcClientAccessServer
```

You can also specify a different CAS.

/CONFIGfile=configfilename

Use the `/CONFIGfile` parameter to specify the name of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file that contains the values for the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Microsoft Exchange Server installation directory is used. If the `/CONFIGfile` parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire `/CONFIGfile` parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

DAGNODE=nodename

Specify the node name that you want to use to back up the databases in an Exchange Server Database Availability Group. With this setting, backups from all Database Availability Group members that are configured to use the DAG node are backed up to a common file space on the IBM Spectrum Protect server. The database copies are managed as a single entity, regardless of which Database Availability Group member they are backed up from. This setting can prevent IBM Spectrum Protect Snapshot for Microsoft Exchange Server from making too many backups of the same database.

/EXCLUDEDUMPster=No|Yes

Use the **/EXCLUDEDUMPster** parameter to specify whether IBM Spectrum Protect Snapshot includes or excludes the mail items in the Recoverable Items folder in mailbox restore operations.

You can specify the following values:

No

Restore the mail items in the Recoverable Items folder to a mailbox restore destination. This option is the default.

Yes

Do not restore the mail items in the Recoverable Items folder to a mailbox restore destination.

If you are restoring the mailbox of George Clark, for example, you can exclude the Recoverable Items folder contents as shown in the following example:

```
tdpexcc restoremailbox "George Clark" /EXCLUDEDUMPster=YES
/USEEXISTINGRDB=NO /KEEPRDB=NO
```

/FCMPTFile=dsmoptfilename

The **/FCMPTFile** parameter specifies the IBM Spectrum Protect Snapshot options file to use.

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/FCMPTFile**, the default value is `dsm.opt`.
- If you specify **/FCMPTFile** but not *dsmoptfilename*, the default is also `dsm.opt`.

/FROMArchive

Use the **/FROMArchive** parameter only if you are restoring an archive mailbox and you specify the mailbox alias of the primary mailbox. If you specify the primary mailbox alias and you do not specify this parameter option, by default, the primary mailbox is restored.

To restore an archive mailbox to another archive mailbox, specify both the **/FROMArchive** and the **/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX, target-mailboxname** parameters. For example:

```
tdpexcc restoremailbox "OriginalArchiveMailboxName" /FROMArchive
/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,"TargetArchiveMailboxName"
```

/KEEPRDB=No|Yes

Use the **/KEEPRDB** parameter to specify whether IBM Spectrum Protect Snapshot retains a recovery database for reuse in mailbox restore operations, or automatically removes the recovery database after mailbox restore operations.

You can specify the following values:

No

Do not retain a recovery database for mailbox restore operations. Remove the recovery database after mailbox restore processing. This option is the default.

Yes

Retain the recovery database for mailbox restore operations.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Microsoft Exchange Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Microsoft Exchange Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpexchange.log"
```

If you do not specify the **/LOGFile** parameter, log records are written to the default log file, `tdpexc.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Microsoft Exchange Server to run operations, use the **/LOGFile** parameter to specify a different log file for each instance that is used. This function directs logging for each instance to a different log file and prevents interspersed log file records.

Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, some days of data are saved. By default, 60 days of log entries are saved. The option No can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify no, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MAILBOXFILTER=ATTACHMENTNAME|ENDDATETIME|FOLDER|MESSAGEBODY|SENDER|STARTDATETIME|SUBJECT|ALLCONTENT

Use the **/MAILBOXFILTER** parameter to specify filters to restrict what mailbox data is restored. When you are restoring to a Unicode .pst file for Exchange Server 2013 or 2016, except for the FOLDER and ALLCONTENT filters, the filters are ignored.

You can specify multiple filters; however, you must specify each filter only one time. For each filter that you specify, a separate **/MAILBOXFILTER** parameter must be used. For example:

```
tdpexcc.exe restoremailbox dchang /MAILBOXFILTER=STARTDATETIME,07/01/2013  
/MAILBOXFILTER=ENDDATETIME,07/31/2013
```

Mailbox data that matches a combination of all filters that are specified is restored. If no filters are specified, by default all data in the mailbox is restored.

Specify one of the following filters when you use this parameter:

ATTACHMENTNAME, *attachmentname-search-text*

Use `/MAILBOXFILTER=attachmentname attachmentname-search-text` to restore only the mailbox messages that contain a match of the specified text within a message attachment name. The match is not case-sensitive. For example, an *attachmentname-search-text* of *Rob* matches the attachment name: *Rob, robert.txt, PROBE, and prObe.pdf*.

Enclose the *attachmentname-search-text* variable in double quotation marks.

The **ATTACHMENTNAME** filter does not match the attachment names of encrypted mailbox messages. If a mailbox message is encrypted, it is skipped by the **ATTACHMENTNAME** filter.

ENDDATETIME, *end-date[,end-time]*

Use `/MAILBOXFILTER=enddatetime, end-date [, end-time]` to restore only the mailbox messages that are sent or received earlier than the specified date and time.

The *end-date* variable is required. Use the same date format for the *end-date* that you selected with the DATEFORMAT option in the IBM Spectrum Protect Snapshot options file.

The *end-time* variable is optional. Use the same time format for the *end-time* variable that you selected with the TIMEFORMAT option in the IBM Spectrum Protect Snapshot options file.

The ENDDATETIME filter date and time must be later than the STARTDATETIME filter date and time. If no time is specified, all messages that are sent or received on that date is restored.

FOLDER, *folder-name*

Use `/MAILBOXFILTER=folder, original-folder-name` to restore only the mailbox messages that are in the specified folder within the original mailbox. The match is not case-sensitive.

Enclose the *original-folder-name* variable in double quotation marks.

- To filter a public folder to restore, ensure that you are restoring the folder to an existing public folder that has the same folder path as the public folder to be restored. If the original public folder is deleted after the time of the backup, re-create the public folder. Specify the full path to the folder. If the full directory path includes spaces, enclose the directory path in double quotation marks, and do not append a backslash character (\) at the end of the directory path.

For example, to restore a folder that is named "SubFolder" under "ParentFolder", specify "ParentFolder/SubFolder" as the folder path. To restore all folders in a parent folder, use *ParentFolder/**.

- To restore a specific folder in an archive mailbox, ensure that you specify the full directory path to the folder.

To restore an archive mailbox to another archive mailbox, you must specify both the `/MAILBOXFILTER=folder, original-folder-name` parameter and the `/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX, target-mailboxname` parameter. For example:

```
tdpexcc restoremailbox "OriginalArchiveMailboxName"  
/MailboxFilter="folder,"folderA" /MAILBOXRESTOREDESTINATION=  
ARCHIVEMAILBOX,"TargetArchiveMailboxName"
```

- To restore the folder of a mailbox to a Unicode .pst file, ensure that you specify the `/MAILBOXFILTER=FOLDER` parameter with the `/`

MAILBOXRESTOREDESTINATION=UNICODEPST parameter. Specify the full directory path to the folder. For example, to restore a folder that is named "SubFolder" under "ParentFolder", specify "ParentFolder/SubFolder" as the folder path. To restore all folders in a parent folder, use ParentFolder/*.

- To restore only the mail items in the Deletions subfolder of the Recoverable Items/ folder, specify the /MAILBOXFILTER=FOLDER parameter with the correct folder value for the target destination.

As shown in the following example, if you are restoring mail items to the original mailbox, specify the Deletions folder.

```
tdpexcc restoremailbox "george clark" /MailboxFilter=folder,
"Deletions"
```

If you are restoring mail items to a Unicode .pst file, specify the full folder path to the Deletions folder.

```
tdpexcc restoremailbox "george clark" /MailboxFilter=folder,
"Recoverable Items/Deletions" /KEEPADB=NO /USEExistingrdb=NO
/MAILBOXRESTOREDESTINATION=UNICODEPST,c:\gclark.pst
```

MESSAGEBODY,messagebody-search-text

Use /MAILBOXFILTER=messagebody, , *messagebody-search-text* to restore only the mailbox messages that contain a match of the specified text within the message body. The match is not case-sensitive. For example, a *messagebody-search-text* of *Rob* matches the message body text: *Rob*, *robert*, *PROBE*, and *prObe*.

Enclose the *messagebody-search-text* variable in double quotation marks.

The MESSAGEBODY filter does not match the message body of encrypted mailbox messages. If a mailbox message is encrypted, it is skipped by the MESSAGEBODY filter.

SENDER,sender-name

Use /MAILBOXFILTER=sender, , *sender-name* to restore only the mailbox messages that are received from the specified message sender.

Enclose the *sender-name* variable in double quotation marks.

STARTDATETIME,start-date[,start-time]

Use /MAILBOXFILTER=startdatetime, *start-date* [, *start-time*] to restore only the mailbox messages that are sent or received after the specified date and time.

The *start-date* variable is required. Use the same date format for the *start-date* that you selected with the DATEFORMAT option in the IBM Spectrum Protect Snapshot options file.

The *start-time* variable is optional. Use the same time format for the *start-time* variable that you selected with the TIMEFORMAT option in the IBM Spectrum Protect Snapshot options file.

The STARTDATETIME filter date and time must be earlier than the ENDDATETIME filter date and time. If no time is specified, all messages that are sent or received on that date is restored.

SUBJECT,subject-search-text

Use /MAILBOXFILTER=subject, , *subject-search-text* to restore only the mailbox messages that contain a match of the specified text within the message subject line. The match is not case-sensitive. For example, a *subject-search-text* of *Rob* matches the subject text: *Rob*, *robert*, *PROBE*, and *prObe*.

Enclose the *subject-search-text* variable in double quotation marks.

ALLCONTENT,allcontent-search-text

Use /MAILBOXFILTER=allcontent, *allcontent-search-text* to restore only the mailbox messages that contain a match of the specified text that is contained within the message sender, the message subject line, or the message body. The match is not case-sensitive. For example, an *allcontent-search-text* of *Rob* matches *Rob*, *robert*, *PROBE*, and *prObe* contained within the attachment name, message sender, the subject line, or the message body.

Enclose the *allcontent-search-text* variable in double quotation marks.

The ALLCONTENT filter does not match the message body of encrypted mailbox messages. If a mailbox message is encrypted, the ALLCONTENT filter matches only text that is contained within the message sender or the subject line.

/MAILBOXORIGLOCATION=server-name,db-name

Use the **/MAILBOXORIGLOCATION** parameter to specify the Exchange Server and the database where the mailbox is at the time of backup.

If you do not specify the **/MAILBOXORIGLOCATION** parameter, the default value is the location (found in the mailbox location history) of the mailbox to restore from, for the backup time specified. If no mailbox location history is available, the default value is the current active location of the mailbox.

server-name

The name of the Exchange Server where the mailbox is at the time of backup.

db-name

The name of the database where the mailbox is at the time of backup.

The **/MAILBOXORIGLOCATION** parameter is only necessary if the mailbox to be restored from is moved or deleted after the time of the backup, and no mailbox location history is available.

A **restoremailbox** operation from a backup that is selected with IBM Spectrum Protect Snapshot for Microsoft Exchange Server before version 6.1 fails if the **/MAILBOXORIGLOCATION** parameter is not specified for mailboxes that meet one or both of the following conditions:

- The mailbox to be restored is moved. (The mailbox is not in the same server and the same database where the mailbox is at the time of backup).
- The mailbox to be restored is deleted and the restore destination is to an alternate mailbox or to a .pst file.

For example:

```
TDPEXCC RESTOREMAILBOX annjones
/MAILBOXORIGLOCATION=serv1,mdbb1
/MAILBOXRESTOREDAt=12/31/2013
/MAILBOXRESTOREDESTination=PST,c:\team99\rcvr.pst
```

/MAILBOXRESTOREDAt=restore-date

Use the **/MAILBOXRESTOREDAt** parameter with or without the **/mailboxrestorettime** parameter to establish a date and time to restore mailbox data from. A mailbox is restored from the earliest backup that is selected after the date and time that is established by the **/MAILBOXRESTOREDAt** and the **/mailboxrestorettime** parameters. Specify the appropriate date in the *restore-date* variable; use the same format that you selected with the DATEFORMAT option in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server options file.

If *restore-date* or *restore-time* values are not specified, no date and time is established. By default the mailbox is restored from the most recent available backup.

If either *restore-date* or *restore-time* is specified, then the mailbox is restored from the earliest backup that is selected after the established restoration date and time. If no backup of the mailbox after the established date and time is found, by default the mailbox will be restored from the most recent available backup.

- If you specify both *restore-date* or *restore-time*, this action establishes the mailbox restoration period.
- If you specify *restore-date* and you do not specify *restore-time*, *restore-time* defaults to a value of 23:59:59. This action establishes the *restore-date* at the specified date.
- If you specify *restore-time* without *restore-date*, then *restore-date* defaults to the current date. This setting establishes the restoration date and time as the current date at the specified *restore-time*.

/MAILBOXRESTORETime=restore-time

Use the **/MAILBOXRESTORETime** parameter with or without the **/MAILBOXRESTOREDate** parameter to establish a date and time to restore a mailbox from. A mailbox is restored from the earliest backup that is selected after the date and time that is established by the **/MAILBOXRESTOREDate** and the **/MAILBOXRESTORETime** parameters. Specify the appropriate time in the *restore-time* variable; use the same format that you selected with the TIMEFORMAT option in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server options file.

If *restore-date* and *restore-time* values are not specified, no date and time is established. By default the mailbox is restored from the most recent available backup.

If either *restore-date* or *restore-time* is specified, the mailbox is restored from the earliest backup that is selected after the established date and time. If no backup of the mailbox after the established date and time is found, by default the mailbox is restored from the most recent available backup.

- If you specify both *restore-date* and *restore-time*, this function establishes the mailbox restoration period.
- If you specify *restore-date* and you do not specify *restore-time*, *restore-time* defaults to a value of 23:59:59. This function establishes the *restore-date* at the specified date.
- If you specify *restore-time* without *restore-date*, the *restore-date* variable defaults to the current date. This function establishes the restoration date and time as the current date at the specified *restore-time*.

/MAILBOXRESTOREDESTINATION=EXCHANGE|PST|UNICODEPST|ARCHIVEMAILBOX

Use the **/mailboxrestoredestination** parameter to specify the destination to restore the mailbox data to.

If you do not specify the **/mailboxrestoredestination** parameter, by default, the EXCHANGE option is used and the **/mailboxrestoredestination** is not required. The default system behavior is to restore mailbox data to the original location in the original active mailbox. When you restore multiple mailboxes with the same **restoremailbox** command, the default system behavior is to restore mailbox data into each original active mailbox.

Mailbox items are merged into the mailbox destination. If a mailbox item exists in the mailbox destination, that item is not restored.

You must specify one of the following values when you use this parameter:

EXCHANGE,[target-mailboxname,target-foldername]

Use the **/mailboxrestoredestination=EXCHANGE** option to restore mailbox messages into a live Exchange Server.

The EXCHANGE option is the default option. If you specify the **/mailboxrestoredestination=EXCHANGE** option without specifying any variables, the result is the same as not specifying the **/mailboxrestoredestination** parameter. The mailbox data is restored to the original location in the original active mailbox.

Use **/mailboxrestoredestination=EXCHANGE, target-mailboxname, target-foldername** to restore mailbox messages into a destination other than the original location in the original active mailbox. The mailbox messages are restored into a subfolder of the specified folder within the target mailbox. The target mailbox can be the original mailbox or an alternate mailbox.

When you restore multiple mailboxes with the same **restoremailbox** command, this option restores the mailbox data into a subfolder (designated by each original mailbox-alias) of the specified target folder in the active mailbox. The folders from the corresponding original mailbox, which contain the restored mailbox messages, are in each subfolder. The specified folder in the target mailbox contains a subfolder that is designated by the original mailbox alias name. Subfolders that contain the restored mailbox messages are in each parent subfolder. These child subfolders have the folder structure of the original mailbox.

target-mailboxname

Specify the target mailbox-alias or the target mailbox-display name. The target mailbox must be an active mailbox.

If the *target-mailboxname* variable includes spaces, enclose the entry in double quotation marks.

To restore a specific public folder to an alternate public folder mailbox, specify both the `/MAILBOXFILTER=folder,original-folder-name` parameter and the `/MAILBOXRESTOREDESTINATION=EXCHANGE,target-publicfolder-mailboxname` parameter. For example:

```
tdpexcc restoremailbox "OriginalPublicFolderMailbox"  
/MailboxFilter=folder,"folderA" /MAILBOXRESTOREDESTINATION=  
EXCHANGE,"TargetPublicFolderMailbox"
```

You can restore a public folder only to an existing public folder on the Exchange server. If the public folder is relocated to an alternate mailbox destination after the time of the backup, ensure that it exists in the alternate mailbox location with the same folder path as the folder to be restored. The restore operation does not automatically re-create the public folder in the destination mailbox.

target-foldername

The *target-foldername* variable specifies the mailbox folder in the target mailbox to restore mailbox messages to.

If you restore a mailbox to a different destination than the original mailbox, the mailbox folders are restored in the destination mailbox under a folder that is named *original-mailbox_mailbox-GUID*. In the process, the Recoverable Items folders are restored.

If you specify the *target-mailboxname* variable and the target mailbox is not the original mailbox, you must specify a folder name. However, when you restore to a mailbox in a target public folder, do not specify a target folder name. A folder name is not required for public folder restore operations.

If the mailbox folder specified by the *target-foldername* variable does not exist in the target mailbox, a folder with the target folder name is created in the target mailbox except for public folder mailboxes.

The target folder contains one subfolder for each original-mailbox that is restored (designated by each original-mailbox alias). The folders from the corresponding original mailbox, which contain the restored mailbox messages, are in each subfolder. If you did not specify the `/mailboxfilter` parameter, the target folder that you specified contains, within the subfolder that is designated by the original mailbox alias, all the folders that are in the mailbox that you are restoring from. If you specified the `/mailboxfilter` parameter, the subfolder within the folder that you specified contains only the folders with messages that match the filter criteria.

If the *target-foldername* variable includes spaces, enclose the entire *target-foldername* variable entry in double quotation marks. For example:

```
/MAILBOXRESTOREDESTINATION=EXCHANGE,Kerry,"temp folder"
```

When you restore multiple mailboxes with the same **restoremailbox** command, and you specify a target folder, each original-mailbox is restored to the target folder in the target mailbox. The target folder contains one subfolder for each original-mailbox that is restored (designated by each original mailbox alias). The folders from the corresponding original mailbox, which contain the restored mailbox messages, are in each subfolder.

For example, this **restoremailbox** operation restores mailboxes "andrew baker" and "sally wood" to the folder "previous_acctmng" in the target mailbox "mary brown":

```
restoremailbox "andrew baker","sally wood"  
/mailboxrestoredest=exchange,"mary brown",previous_acctmng
```

The restored mailbox messages are placed in folders that are copied from the original mailboxes that use the following folder structure:

```
mary brown (target mailbox)
  >-previous_acctmngr (specified folder)
  |
  | >-abaker (original-mailbox1 alias)
  | | >-Inbox (restored folder from mailbox1)
  | | >-Outbox (restored folder from mailbox1)
  | | >-My Accts (restored folder from mailbox1)
  | | >-swood (original-mailbox2 alias)
  | | >-Inbox (restored folder from mailbox2)
  | | >-Outbox (restored folder from mailbox2)
  | | >-New Accts (restored folder from mailbox2)
```

PST,non-Unicode-pst-filename-path

Use `/mailboxrestoredestination=PST,non-Unicode-pst-filename-path` to restore mailbox data to an Exchange Server personal folders (.pst) file. The mailbox data that is restored is in non-Unicode format.

You can include the `non-Unicode-pst-filename-path` variable to specify the destination where the **restoremailbox** operation writes the .pst file. The `non-Unicode-pst-filename-path` can be either a fully qualified path to a .pst file or a directory path. If you do not specify a path, the .pst file is written to the current directory.

- You can specify `non-Unicode-pst-filename-path` as a fully qualified path to a .pst file to restore all mail to that .pst file.

```
TDPEXCC RESTOREMAILBOX gclark
  /mailboxrestoredestination=PST,c:\mb\dept54\vpo.pst
```

Note: The .pst directory must exist before you use the **restoremailbox** command. The .pst file is created if it does not exist.

If you are restoring more than one mailbox and you specify a fully qualified path to a .pst file, all the mailbox data is restored to the one .pst file specified. Inside the .pst file, the parent-level folder name is the mailbox-alias-name, followed by the rest of the mailbox folders.

- You can specify `non-Unicode-pst-filename-path` as a directory path to have IBM Spectrum Protect Snapshot for Microsoft Exchange Server create a .pst file by using the mailbox-alias-name of the mailbox that is being restored, and store the .pst file in the specified directory. For example, the .pst file name of the restored mailbox "George Clark" (gclark) is `gclark.pst`.

```
TDPEXCC RESTOREMAILBOX "george clark"
  /mailboxrestoredestination=PST,c:\mb\dept54\
```

The .pst directory must exist before you use the **restoremailbox** command. If the .pst file does not exist, the file is created.

If you restore multiple mailboxes with the same **restoremailbox** command, and you specify a directory path, each mailbox is restored into a separate .pst file. For example, if mailboxes John (john1), John Oblong (oblong), and Barney Olef (barneyo) are restored and the specified directory path is `c:\finance`, all mailboxes are restored into the `c:\finance` directory as shown:

```
c:\finance\john1.pst
c:\finance\oblong.pst
c:\finance\barneyo.pst
```

The .pst directory must exist before you use the **restoremailbox** command. The mailbox data that is restored by using `/mailboxrestoredestination=PST,non-Unicode-pst-filename-path` must be less than 2 GB.

If the *non-Unicode-pst-filename-path* variable includes spaces, enclose the entire *non-Unicode-pst-filename-path* variable entry in double quotation marks and do not append a backslash character (\) at the end of folder path. For example:

```
TDPEXCC RESTOREMAILBOX "george clark"  
/mailboxrestoredestination=PST,"c:\mb\dept54\access group"
```

UNICODEPST,Unicode-pst-filename-path

Use `/mailboxrestoredestination=UNICODEPST,Unicode-pst-filename-path` to restore mailbox data to an Exchange Server personal folders (.pst) file. The mailbox data that is restored is in Unicode format.

You can include the *Unicode-pst-filename-path* variable to specify where the **restoremailbox** operation locates the .pst file. The *Unicode-pst-filename-path* can be either a fully qualified UNC path to a .pst file or a directory path. If you do not specify a path, the .pst file is written to the current directory. If you specify a non-UNC path (such as `c:\dir\mailbox.pst`), IBM Spectrum Protect Snapshot for Microsoft Exchange Server tries to convert it to a UNC path for you, but it might not work for custom UNC paths or shares.

- To restore all mail to a .pst file, specify *Unicode-pst-filename-path* as a fully qualified path to the .pst file.

```
TDPEXCC RESTOREMAILBOX gclark  
/mailboxrestoredestination=UNICODEPST,c:\mb\dept54\vpo.pst
```

If the *Unicode-pst-filename-path* variable includes spaces, enclose the entire *Unicode-pst-filename-path* variable entry in double quotation marks and do not append a backslash character (\) at the end of folder path. For example:

```
TDPEXCC RESTOREMAILBOX "george clark"  
/mailboxrestoredestination=UNICODEPST,"c:\mb\dept54\access group"
```

The .pst directory must exist before you issue the **restoremailbox** command. If the .pst file does not exist, the file is created.

If you are restoring more than one mailbox and you specify a fully qualified path to a .pst file, all the mailbox data is restored to the .pst file that you specify. The parent-level folder name in the .pst file is the mailbox-alias-name. The remaining mailbox folders follow the parent-level folder.

- Specify *Unicode-pst-filename-path* as a directory path if you want IBM Spectrum Protect Snapshot for Microsoft Exchange Server to create a .pst file by using the mailbox-alias-name of the mailbox that is being restored, and to store the .pst file in the specified directory. For example, the .pst file name of the restored mailbox "George Clark" (gclark) is `gclark.pst`.

```
TDPEXCC RESTOREMAILBOX "george clark"  
/mailboxrestoredestination=UNICODEPST,c:\mb\dept54
```

The .pst directory must exist before you issue the **restoremailbox** command. If the .pst file does not exist, the file is created.

If you restore multiple mailboxes with the same **restoremailbox** command, and you specify a directory path, each mailbox is restored into a separate .pst file. For example, if mailboxes John (john1), John Oblong (oblong), and Barney Olef (barneyo) are restored and the specified directory path is `c:\finance`, all mailboxes are restored into the `c:\finance` directory as shown:

```
c:\finance\john1.pst  
c:\finance\oblong.pst  
c:\finance\barneyo.pst
```

- To restore only the mail items in the Deletions subfolder of the Recoverable Items/ folder, specify the /MAILBOXFILTER=FOLDER parameter with the correct folder value for the target destination.

As shown in the following example, if you are restoring mail items to a Unicode .pst file, specify the full folder path to the Deletions folder.

```
tdpexcc restoremailbox "george clark" /MailboxFilter=folder,
"Recoverable Items/Deletions" /KEEPADB=NO /USEExistingrdb=NO
/MAILBOXRESTOREDESTINATION=UNICODEPST,c:\gclark.pst
```

ARCHIVEMAILBOX,[target-mailboxname,target-foldername]

Use **/MAILBOXRESTOREDESTINATION** with the ARCHIVEMAILBOX and **/FROMARCHIVE** parameters to restore archive mailbox messages to its original archive mailbox or to an alternate archive mailbox.

Use **/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX, target-mailboxname** to specify the archive mailbox destination that you want to restore to. You can also specify a target folder name in the archive mailbox.

To restore an archive mailbox into a specific folder of an archive mailbox, specify both the **/FROMArchive** parameter and the **/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX, target-mailboxname, target-foldername** parameters. For example:

```
tdpexcc restoremailbox "OriginalFolderName" /FROMArchive
/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX,"TargetFolderName"
"folderA"
```

If you specify the **/MAILBOXRESTOREDESTINATION=ARCHIVEMAILBOX** parameter without specifying a target mailbox destination, the mailbox messages are restored to the original location in the original archive mailbox.

/MAILBOXRestoreunread=Yes|No

Use the **/MAILBOXRestoreunread** parameter to specify whether IBM Spectrum Protect Snapshot for Microsoft Exchange Server marks restored mail messages as unread.

You can specify the following values:

Yes

Mark restored mail messages as unread. This option is the default.

No

Do not mark restored mail messages as unread.

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

A backup is mounted as read-only, which results in a file copy of the Microsoft Exchange database file to the RDB to complete the mailbox restore operation.

Yes

A backup is mounted as read/write to do the mailbox restore operation. The backup is mounted on the directory you specify with the **/MOUNTPath** parameter. If a **/MOUNTPath** value is not specified, a temporary directory is used (system environment TEMP variable).

Note: When you specify the **/MOUNTRW** parameter for mailbox restore operations, the **/USEEXISTINGRDB** and **/KEEPRDB** parameters also apply.

- If both **/MOUNTRW** and **/USEEXISTINGRDB** are set to Yes and a recovery database (RDB) exists on the system, the existing RDB is used for mailbox restore operations and **/MOUNTRW** is ignored.
- If **/KEEPRDB** is specified, the snapshot RDB remains mounted on the system after the mailbox restore operation is complete (you must remove the snapshot RDB manually later). If you also specified the **/MOUNTRW** parameter, you must unmount the RDB by using the **unmount backup** command or the Windows Powershell cmdlet **Dismount-DpExcBackup**.

/Quiet

This parameter prevents the display of status information but does not affect the level of information that is written to the activity log.

/TEMPDBRESTorepath=*path-name*

Use the **/TEMPDBRESTorepath** parameter to specify the default temporary path to use when you restore mailbox database files.

If you do not specify the **/TEMPDBRESTorepath** parameter, the default value is the value that is specified by the **/TEMPDBRESTorepath** option in the IBM Spectrum Protect Snapshot configuration file. The default IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file is `tdpexc.cfg`. If the **/TEMPDBRESTorepath** value does not exist in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file, the TEMP environment variable value is used.

If the *path-name* variable includes spaces, enclose the entire **/TEMPDBRESTorepath** parameter entry in double quotation marks. For example:

```
TDPEXC RESTOREMAILBOX richgreene  
/tempdbrestorepath="h:\Exchange Restore Directory"
```

- Do not specify a value of **/TEMPDBRESTorepath** that is the same value as the location of the active database. If the value is the same, the database might become corrupted.
- Choose a temporary database-restore location that has enough space to hold the entire restore for the database.

For better performance, the current active-transaction logger is to be on a different physical device from the paths that are specified by the values of the **/TEMPDBRESTorepath** parameter and the **/TEMPDBRESTorepath** parameter. The paths that are specified by the values of the **/TEMPDBRESTorepath** parameter and the **/TEMPDBRESTorepath** parameter can be on the same or separate physical devices from each other.

Do not specify double-byte characters (DBCS) within the temporary database-restore path.

/TEMPLOGRESTorepath=*path-name*

Use the **/TEMPLOGRESTorepath** parameter to specify the default temporary path to use when you restore logs and patch files.

If you do not specify the **/TEMPLOGRESTorepath** parameter, the default value is the value that is specified by the **/TEMPLOGRESTorepath** option in the IBM Spectrum Protect Snapshot configuration file. The default IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file is `tdpexc.cfg`. If you do not specify the **/TEMPLOGRESTorepath** parameter and the **/TEMPLOGRESTorepath** value does not exist in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file, the TEMP environment variable value is used.

- Do not specify a value of **/TEMPLOGRESTorepath** that is the same value as the current location for the database that is used for recovery. If the value is the same, the database might become corrupted.
- Choose a temporary log-restore location that has enough space to hold all the log and patch files.

For better performance, the current active-transaction logger is to be on a different physical device from the paths that are specified by the values of the **/TEMPLOGRESTorepath** parameter and the **/TEMPLOGRESTorepath** parameter. The paths that are specified by the values of the **/TEMPLOGRESTorepath** parameter and the **/TEMPLOGRESTorepath** parameter can be on the same or separate physical devices from each other.

TEMPLOGRESTorepath parameter and the **/TEMPLOGRESTorepath** parameter can be on the same or separate physical devices from each other.

Do not specify double-byte characters (DBCS) within the temporary log-restore path.

/TEMPMAILBOXAlias=tempmailbox-alias

Use the **/TEMPMAILBOXAlias** parameter to specify the mailbox-alias of a temporary mailbox to use.

If you do not specify the **/TEMPMAILBOXAlias** parameter, the default value is the value that is specified by the **/TEMPMAILBOXAlias** option in the IBM Spectrum Protect Snapshot configuration file. The default IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file is `tdpexc.cfg`. If the **/TEMPMAILBOXAlias** value does not exist in the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file, the mailbox of the currently logged on user is used as the temporary mailbox.

Specify the following value when you use this parameter:

tempmailbox-alias

Specify the mailbox-alias of the temporary mailbox to use for recovery of mailboxes that are deleted or re-created after the time of the backup you are restoring from.

Ensure that the temporary mailbox is active and has enough storage capacity to accommodate all items of the mailboxes that are being restored.

If the *tempmailbox-alias* variable includes spaces, enclose the entry in double quotation marks.

/USEEXISTINGRDB=Yes|No

Use the **/USEEXISTINGRDB** parameter to specify whether IBM Spectrum Protect Snapshot for Microsoft Exchange Server restores mailboxes from an existing recovery database, or automatically removes an existing recovery database during mailbox restore operations.

You can specify the following values:

Yes

Use an existing recovery database for mailbox restore operations. This option is the default.

No

Do not use an existing recovery database for mailbox restore operations. Remove the recovery database during mailbox restore processing.

Examples: restoremailbox command

You can combine the use of the **/KEEPRDB** and **/USEEXISTINGRDB** parameter options with the **restoremailbox** command.

Example: Use an existing recovery database for mailbox operations

Use an existing recovery database for restore mailbox operations so that you do not have to restore the recovery database again.

```
tdpexcc restoremailbox <MB> /USEEXISTINGRDB=Yes
```

Example: Retain a recovery database for mailbox operations

Retain a recovery database after a mailbox restore operation so that you can use the recovery database for other restore operations.

```
tdpexcc restoremailbox <MB> /KEEPRDB=YES
```

Example: Retain a recovery database for multiple mailbox restore operations, and then remove it

Because you restore multiple mailboxes at different times, you want to retain the recovery database after the first mailbox restore operation and use it for subsequent restore operations. When you restore the final mailbox, you remove the recovery database.

```
tdpexcc restoremailbox <MB_1> /KEEPRDB=YES
```

```
tdpexcc restoremailbox <MB_2> /USEEXISTINGRDB=YES
```

```
tdpexcc restoremailbox <MB_n> /KEEPRDB=NO
```

Example: Restore multiple mailboxes simultaneously

Simultaneously restore multiple mailboxes and ensure that the recovery database is automatically removed after each mailbox is restored.

```
tdpexcc restoremailbox <MB_1>,<MB_2> /KEEPRDB=NO
```

Example: Restore multiple mailboxes from an existing recovery database

Simultaneously restore multiple mailboxes from an existing recovery database.

Tip: Mailboxes that are not in the recovery database are bypassed during restore processing, and are indicated in the console output.

Restore the remaining mailboxes that are not in the recovery database.

```
tdpexcc restoremailbox <MB_1>,<MB_2>...<MB_n> /USEEXISTINGRDB=YES  
/KEEPRDB=NO
```

```
tdpexcc restoremailbox <MB_1>,<MB_2>...<MB_n> /USEEXISTINGRDB=NO  
/KEEPRDB=NO
```

Set command

Use the **set** command to set the IBM Spectrum Protect Snapshot for Exchange Server configuration parameters in a configuration file.

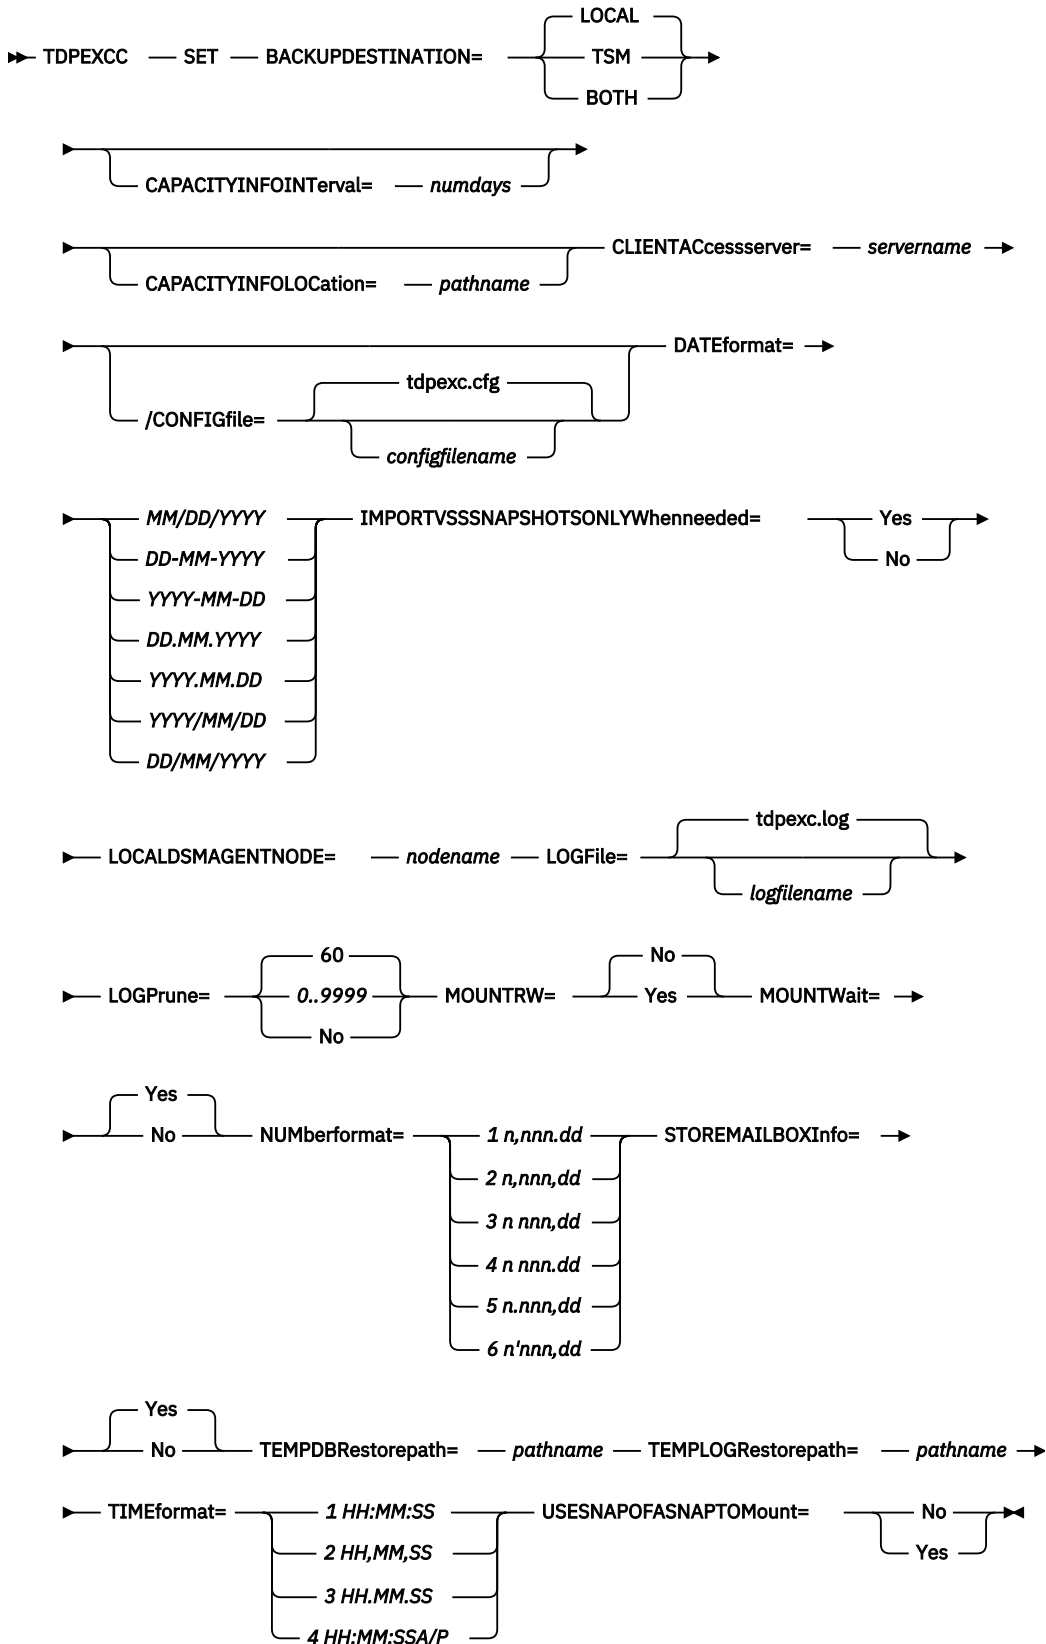
The values that you set are saved in an IBM Spectrum Protect Snapshot for Exchange Server configuration file. The default file is `tdpexc.cfg`. Configuration values can also be set in the **Data Protection Properties** window in Microsoft Management Console (MMC).

For command invocations other than this command, the value of a configuration parameter that is specified in a command overrides the value of the configuration parameter that is specified in the IBM Spectrum Protect Snapshot for Exchange Server configuration file. If, when you use this command, you do not override a value for the configuration file parameter, the values in the default configuration file are used.

Set syntax

Use the **set** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command



Set positional parameters

Positional parameters immediately follow the **set** command and precede the optional parameters.

The following positional parameters specify the values in the IBM Spectrum Protect Snapshot for Exchange Server configuration file. You can set only one value for each **tdpexcc set** command run:

BACKUPDESTination=TSM|LOCAL|BOTH

Use the **BACKUPDESTination** positional parameter to specify the storage location for your backup. You can specify:

TSM

The backup is stored on IBM Spectrum Protect server storage only. This option is the default.

LOCAL

The backup is stored on local shadow volumes only.

BOTH

The backup is stored on both IBM Spectrum Protect server storage and local shadow volumes.

CAPACITYINFOInterval=numdays

Use the **CAPACITYINFOInterval** positional parameter to specify how often you want the capacity metrics report to be generated. The report, in an XML file format, is generated automatically at the end of a backup operation. The valid value range is 1 - 365 and the default value is 7 days, which means the report is generated once every 7 days.

CAPACITYINFOLOcation=pathname

Use the **CAPACITYINFOLOcation** positional parameter to specify the location where the capacity metrics report is to be created. If you do not specify a location, the report is not generated.

CLIENTACcessserver=servername

The *servername* variable refers to the name of the server you use to access the client.

DATEformat=dateformatnum

Use the **DATEformat** positional parameter to select the format you want to use to display dates.

The *dateformatnum* variable displays the date in one of the following formats. Select the format number that corresponds to the format you want to use.

1

(Default) *MM/DD/YYYY*

2

DD-MM-YYYY

3

YYYY-MM-DD

4

DD.MM.YYYY

5

YYYY.MM.DD

6

YYYY/MM/DD

7

DD/MM/YYYY

Changes to the value of the **DATEformat** parameter can result in an undesired pruning of the IBM Spectrum Protect Snapshot for Exchange Server log file (*tdpexc.log* by default). You can avoid losing existing log file data by doing one of the following choices:

- After you change the value of the **DATEformat** parameter, make a copy of the existing log file before you run IBM Spectrum Protect Snapshot for Exchange Server.
- Specify a new log file with the **LOGFile** parameter.

IMPORTVSSSNAPSHOTSONLYWhenneeded

Use the **/IMPORTVSSSNAPSHOTSONLYWhenneeded** parameter to specify whether IBM Spectrum Protect Snapshot automatically imports VSS snapshots to the Windows system where the snapshots are created.

Specify one of the following values:

Yes

Import VSS snapshots to the Windows system where the snapshots are created. The option is the default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required. This option is the default.

Tip: For instant restore processing on IBM and non-IBM storage systems, you must specify the Yes option to enable the storage system to create transportable snapshots during backups.

No

Do not create transportable VSS snapshots during backup processing, and do not automatically import the snapshot to storage systems after the backup is completed.

LOCALDSMAgentnode=nodename

Specify the node name of the local system that runs the VSS backups. This positional parameter must be specified for VSS operations to run.

LOGFile=logfilename

Use the **LOGFile** positional parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server. The IBM Spectrum Protect Snapshot for Exchange Server activity log records significant events, such as completed commands and error messages.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is assigned to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. In the configuration file, the default value for the **LOGPrune** is that specified by the **logprune** configurable option. The default value is *60*, which means 60 days of log entries are saved. The option **No** can be specified to disable log pruning.

Regardless of the option that is set in the configuration file for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify the **LOGPrune** parameter, that value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **LOGFile** parameter or **logfile** setting.

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

Perform a read-only mount operation.

Yes

Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPMount** parameter in the configuration file.

- If **USESNAPOFASNAPMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
- If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which require IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which require IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate more target volumes on your storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

MOUNTWait=Yes|No

Use the **MOUNTWait** positional parameter to specify whether IBM Spectrum Protect Snapshot for Exchange Server waits for removable media to mount (such as tapes or DVDs) or to stop the current operation. This situation occurs when the IBM Spectrum Protect server is configured to store backup data on removable media and waits for a required storage volume to be mounted.

Specify **Yes** for IBM Spectrum Protect Snapshot for Exchange Server to wait until all initial volumes of any required removable media are made available to the IBM Spectrum Protect server before you complete the command.

Specify **No** for IBM Spectrum Protect Snapshot for Exchange Server to end the command (if removable media are required). An error message is displayed.

NUMBERformat=fmtnum

Use the **NUMBERformat** positional parameter to specify the format you want to use to display numbers.

The *fmtnum* variable displays numbers by using one of the following formats. Select the format number that corresponds to the format you want to use.

1

(Default) *n,nnn.dd*

- 2 *n,nnn,dd*
- 3 *n nnn,dd*
- 4 *n nnn.dd*
- 5 *n.nnn,dd*
- 6 *n'nnn,dd*

STOREMAILBOXInfo=Yes|No

The **STOREMAILBOXInfo** parameter is used to track mailbox history for moved and deleted mailboxes. By default, this parameter is set to Yes. If you do not plan to use mailbox restore, you can set this option to No. When the option is set to No, IBM Spectrum Protect Snapshot for Exchange Server does not back up the mailbox history.

In large or geographically dispersed domains, more time is required to complete the backup mailbox history task. In this scenario, you can reduce the amount of time that is required to complete the backup mailbox history task by setting the option for **STOREMAILBOXInfo** to No. When you set the option for **STOREMAILBOXInfo** to No, mailboxes that are not moved or are not deleted can be restored normally. Moved and deleted mailbox restores can use the **MAILBOXORIGLOCATION** parameter (of the **restoremailbox** command) to specify the mailbox location at the time of the backup.

TEMPDBRESTorepath=pathname

For mailbox restore operations, use the **TEMPDBRESTorepath** positional parameter to specify the default temporary path to use when you restore mailbox database files.

If you do not enter a path, the default value is the value of the TEMP environment variable.

If the path name includes spaces, you must enclose the entire **TEMPDBRESTorepath** positional parameter entry in double quotation marks. For example:

```
TDPEXCC SET TEMPDBRESTorepath="h:\Exchange Restore Directory"
```

Do not specify a value of **TEMPDBRESTorepath** that is the same value as the location of the active database. If the value is the same, the database might become corrupted. Choose a temporary database-restore location that has enough space to hold the entire restore.

For better performance, the current active-transaction logger must be on a different physical device from the paths that are specified by the values of the **TEMPDBRESTorepath** parameter setting and the **TEMPDBRESTorepath** parameter setting. The paths that are specified by the values of the **TEMPDBRESTorepath** parameter setting and the **TEMPDBRESTorepath** parameter setting can be on the same or separate physical devices from each other.

Do not specify double-byte characters (DBCS) within the temporary database-restore path.

TIMEformat=formatnumber

Use the **TIMEformat** positional parameter to specify the format in which you want system time that is displayed.

The *formatnumber* variable displays time in one of the following formats. Select the format number that corresponds to the format you want to use.

- 1 (Default) *HH:MM:SS*
- 2 *HH,MM,SS*
- 3 *HH.MM.SS*

USESNAPOFASNAPTOmount=Yes|No

During mount operations, you can specify that you want to do a read/write mount by setting **/MOUNTRW=Yes**. When you set the **/MOUNTRW=Yes**, the **USESNAPOFASNAPTOmount** parameter applies and you can further specify whether you want to mount an existing backup or to create a snapshot of an existing backup. You can only set the **USESNAPOFASNAPTOmount** parameter in your configuration file.

- If **USESNAPOFASNAPTOmount** is set to No, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected on the **VSS Options** properties page. After mounting, the original COPY backup can be modified and so can no longer be used as a restore point for future database restore operations.
- If **USESNAPOFASNAPTOmount** is set to Yes, the **Mount read/write (without modifying backup)** check box is selected on the **VSS Options** properties page. This option is only available for SAN Volume Controller (SVC) devices.

Important: You can set **USESNAPOFASNAPTOmount=Yes** only for SAN Volume Controller (SVC) devices with IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Also, you must allocate more target volumes on your SVC storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume matching the size of the volume to be mounted is needed for each concurrent read/write mount of that volume.

Set optional parameters

Optional parameters follow the **set** command and positional parameters.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name of the IBM Spectrum Protect Snapshot for Microsoft Exchange Server configuration file in which these values are set.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for Microsoft Exchange Server installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

Perform a read-only mount operation.

Yes

Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPTOmount** parameter in the configuration file.

- If **USESNAPOFASNAPTOmount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in

future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).

- If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which require IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which require IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate more target volumes on your storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

Set example

The **tdpexcc set localdsmagentnode=mean** command sets the node *mean* as the node name of the local system that processes the backups.

Specify the node name of the local system that processes the VSS backups. When the command completes, the following message is displayed:

```
FMX5054I The preference has been set successfully.
```

Unmount backup command

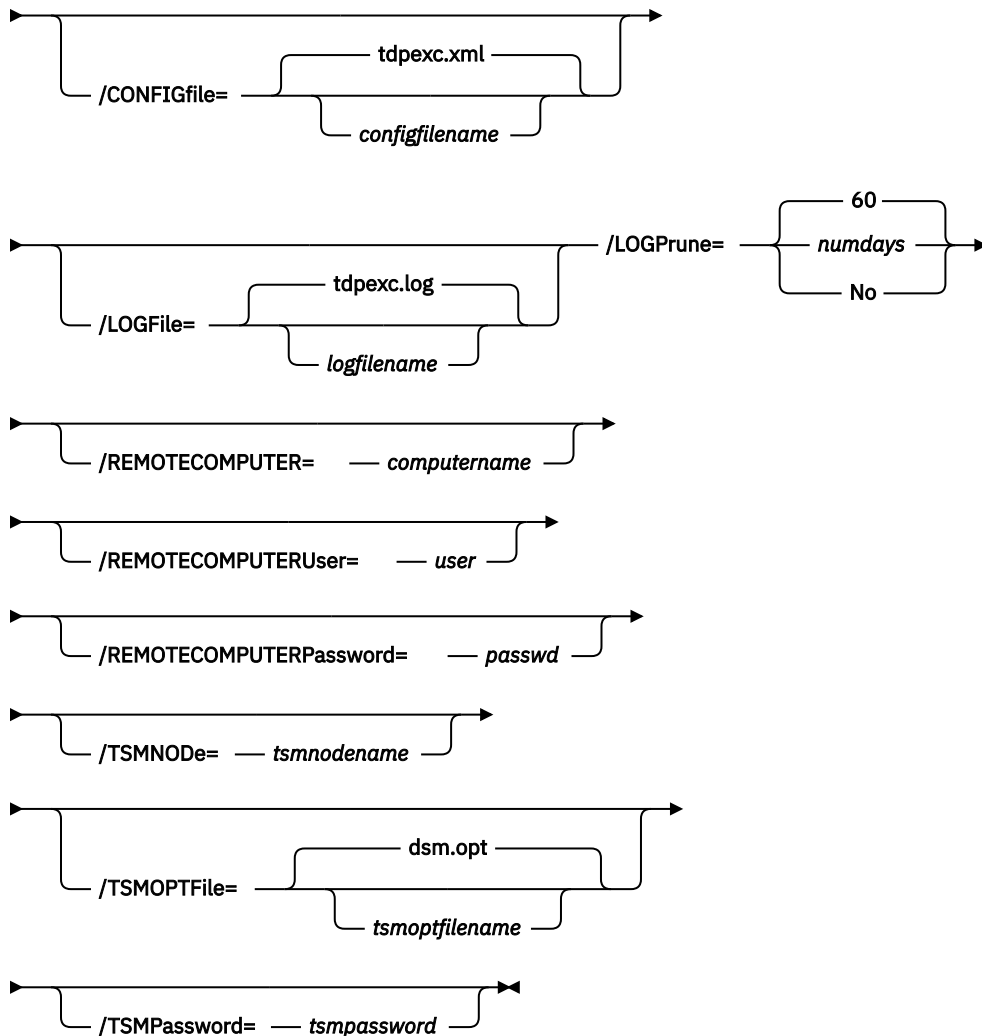
Use the **unmount backup** command to unmount backups that were previously mounted, and are managed by IBM Spectrum Protect Snapshot for Exchange Server.

Unmount backup syntax

Use the **unmount backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPEXCC command

► TDPEXCC — UNMOUNT BACKUP — *mount point root directory* ►



Unmount backup positional parameter

The positional parameter immediately follows the **unmount backup** command and precedes the optional parameters.

mount points root directory

Absolute path to the directory where the snapshots are displayed as mount point directories.

Unmount backup optional parameters

Optional parameters follow the **unmount backup** command and positional parameters.

/CONFIGfile=configfilename

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the configuration file that contains the values to use for an **unmount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpexc.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpexc.cfg"
```

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Exchange Server. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Exchange Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpexc.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `tdpexc.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/REMOTECOMPUTER=computername

Enter the computer name or IP address of the remote system where the backup was created.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/TSMOPTFile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS** GENERATE in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS** GENERATE is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS** PROMPT is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Unmount backup example

This output example provides a sample of the text, messages, and process status that displays when you use the **unmount backup** command.

For a local backup, enter the following command:

```
tdpexcc unmount backup C:\mount-points-root-dir
```

For a remote backup, enter the following command:

```
tdpexcc unmount backup C:\mount-points-root-dir /remotecomputer=computer-name  
/remotecomputeruser=userID /remotecomputerpassword=user password
```

Command-line overview: IBM Spectrum Protect Snapshot for SQL Server

The name of the IBM Spectrum Protect Snapshot for SQL Server command-line interface is `tdpsqlc.exe`. If you installed the **TDPSQL** package, or you configured the Microsoft SQL Server in Microsoft Management Console (MMC), the program is located (by default) in the IBM Spectrum Protect Snapshot installation directory (`C:\Program Files\Tivoli\tsm\TDPSQL`).


Command-line parameter characteristics

The command-line parameters have the following characteristics:

- Positional parameters do not include a leading slash (/) or dash (-).
- Optional parameters can display in any order after the required parameters.
- Optional parameters begin with a forward slash (/) or a dash (-).
- Minimum abbreviations for keywords are indicated in uppercase text.
- Some keyword parameters require a value.
- For those keyword parameters that require a value, the value is separated from the keyword with an equal sign (=).
- If a parameter requires more than one value after the equal sign, the values are separated with commas.
- Each parameter is separated from the others by using spaces.
- If a parameter value includes spaces, the value must be enclosed in double quotation marks.
- A positional parameter can display only once per command invocation.

Where repeatable syntax displays, separate multiple values with commas as indicated in the following example:

TDPSQLC command



To select all instances on the server of database names or file names, specify the asterisk (*) wildcard character following the command.

Command-line interface help

Issue the `tdpsqlc ?` or `tdpsqlc help` command to display help for the command-line interface. You can see more specific help for commands by entering a command like the following example: **tdpsqlc help backup**, where **backup** is an example of a command.

Related tasks

Protecting SQL Server data

With IBM Spectrum Protect Snapshot for SQL Server, you can back up and restore Microsoft SQL Server databases in a stand-alone configuration.

Backup command

Use the **backup** command to back up all or part of one or more SQL databases from the SQL Server to IBM Spectrum Protect Snapshot.

You can enter the asterisk (*) wildcard character to back up all databases. You can specify more than one database for multiple database and transaction log backups.


When you use the **backup** command, remember the following facts:

- You cannot back up or restore the tempdb database because this database is created by the SQL Server each time the server is started.
- The user ID that is used by IBM Spectrum Protect Snapshot to log on to the SQL Server must have the SQL Server SYSADMIN fixed server role.
- You can use the TRANSACT-SQL database consistency checker statement DBCC CHECKDB ('DBNAME') to verify the integrity of the SQL databases before you back them up.


Backup syntax

Use the **backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



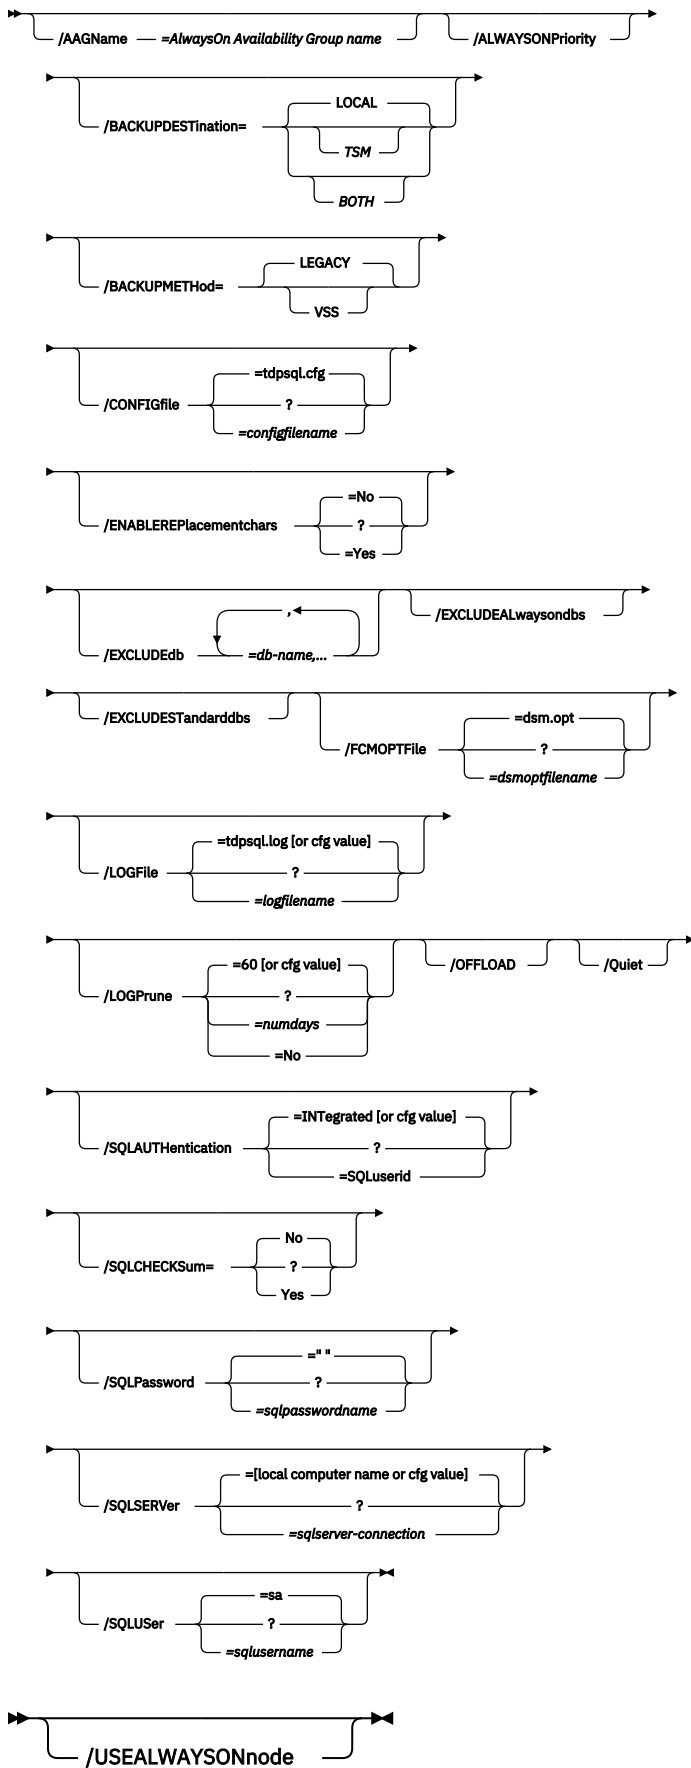
Backup optional parameters




Notes:

¹ For the optional parameters, the **/BACKUPMETHod=** is only valid when using the **full** or **copyfull** positional parameters. The **full** and **copyfull** backups can be performed using VSS or legacy operations. The **/BACKUPMETHod=** parameter is used to choose between the options. The **log**, **diff**, **file**, and **group** backups can be performed only when using legacy operations. You cannot specify the **/BACKUPMETHod=** parameter with these types of backups because only legacy backups are viable.

² The **/BACKUPDESTination** parameter is valid only when using the **full** or **copyfull** positional parameters. The **full** and **copyfull** backups can be saved to local storage, TSM server storage, or both. The **/BACKUPDESTination** parameter is used to choose among the options.



Log Options



Backup positional parameters

Positional parameters immediately follow the **backup** command and precede the optional parameters.

The following positional parameters specify the object to back up:

* | *dbname*

Back up all databases. Use caution when you specify the wildcard character (*) as Microsoft warns not to back up more than a few dozen databases in a single command because of SQL Server limitations.

dbname

Back up the specified database. Multiple entries are separated by commas. If separated by commas, ensure that there is no space between the comma and the database name. If any database name contains blanks, enclose the database name in double quotation marks.

The following positional parameter specifies the type of backup to run:

FULL

A **full** VSS database backup contains all of the contents of a SQL Server database, such as database files, log files, full-text index files, and FILESTREAM files (SQL Server 2008 or later versions).

COPYFull

A copy-only full backup contains a copy-only version of a full backup. These backups are considered out of the regular sequence of conventional SQL Server backups. The backups do not affect the transaction logs or any sequence of backups, such as differential backups or full backups. Use this option to create copy-only full backups periodically for long-term retention without affecting existing backup schedules or retention policies for disaster recovery.

Backup optional parameters

Optional parameters follow the **backup** command and positional parameters.

/AAGName=AlwaysOn Availability Group name

When you back up a database list or all databases by specifying the asterisk (*) wildcard character, and you specify the **/AAGName** parameter, only databases from the availability group that you specify are backed up.

/ALWAYSONPriority

Use this parameter to specify that a local availability database is backed up only if it has the highest backup priority among the availability replicas that are working properly on SQL Server 2012 and later versions. You can use this parameter at the command-line interface or as part of a scheduled backup.

/BACKUPDESTination= LOCAL|TSM|BOTH

Use the **/BACKUPDESTination** parameter to specify the location where the backup is stored.

You can specify:

IBM Spectrum Protect

The backup is stored on IBM Spectrum Protect server storage only. This option is the default.

LOCAL

The backup is stored on local shadow volumes only. This option is only valid when the **/BACKUPMETHod** parameter specifies VSS.

BOTH

The backup is stored on IBM Spectrum Protect server storage and local shadow volumes. This option is valid only when the **/BACKUPMETHOD** parameter specifies VSS.

The **/BACKUPDESTINATION** parameter is valid only when the **full** or **copyfull** positional parameters are used. The **full** and **copyfull** backups can be saved to IBM Spectrum Protect server storage, local storage, or both. The **/BACKUPDESTINATION** parameter is used to choose among options. The **log**, **diff**, **file**, and **group** backups can be stored only to IBM Spectrum Protect server storage. In this scenario, you cannot specify the **/BACKUPDESTINATION** parameter because IBM Spectrum Protect is the only viable option.

/BACKUPMETHOD=LEGACY|VSS

Use the **/BACKUPMETHOD** parameter to specify the manner in which the backup is completed.

You can specify:

LEGACY

The backup is completed with the legacy API. This backup is the SQL streaming backup and restore API as used in previous versions of IBM Spectrum Protect Snapshot for SQL. This option is the default value.

VSS

The backup is completed with VSS.

The **/BACKUPMETHOD** parameter is valid only when the **full** or **copyfull** positional parameters are used. The **full** and **copyfull** backups can be completed by using VSS or legacy operations. The **/BACKUPMETHOD** parameter is used to choose between the options. The **log**, **diff**, **file**, and **group** backups can be completed only by using legacy operations. In this scenario, you cannot specify the **/BACKUPMETHOD** parameter because the legacy method is the only viable option.

/CONFIGfile=configfilename

The **/CONFIGfile** parameter specifies the name of the IBM Spectrum Protect Snapshot configuration file. The configuration file contains the values for the IBM Spectrum Protect Snapshot configurable options. When you use this parameter, review the following information:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where IBM Spectrum Protect Snapshot is installed.
- If *configfilename* includes spaces, place the space character in double quotation marks.
- If you do not specify **/CONFIGfile**, the default value is `tdpsql.cfg`.

/ENABLEREplacementchars=No|Yes

The **/ENABLEREplacementchars** parameter enables SQL Server databases that have backslash (\) or colon (:) characters in the database name to be backed up. The maximum length of the database name is 128 characters. This parameter applies only to IBM Spectrum Protect Snapshot for SQL Server version 7.1.1 and later versions.

You can specify the following values:

Yes

Enable IBM Spectrum Protect Snapshot for SQL Server to process backslash (\) or colon (:) characters in a database name, and back up the database to IBM Spectrum Protect. This value is the default.

No

Prevent database backups to IBM Spectrum Protect if a user-defined string is substituted for a backslash (\) or colon (:) character in the database name.

/EXCLUDEdb=dblist

The **/EXCLUDEdb** parameter specifies the name of the databases to exclude from the backup operation.

/EXCLUDEAlwaysondbs

Use this parameter to exclude all AlwaysOn Availability Databases from the backup operation. If you want to exclude specific databases, use the **/EXCLUDEdb** parameter.

/EXCLUDEStandarddbs

Use this parameter to exclude all standard databases from the backup operation. If you want to exclude specific databases, use the **/EXCLUDEdb** parameter.

/FCMOPTFile=*dsmoptfilename*

The **/FCMOPTFile** parameter specifies the IBM Spectrum Protect Snapshot options file to use.

Considerations:

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/FCMOPTFile**, the default value is *dsm.opt*.
- If you specify **/FCMOPTFile** but not *dsmoptfilename*, the default is also *dsm.opt*.

/LOGFile=*logfile*

The **/LOGFile** parameter specifies the name of the activity log that is generated by IBM Spectrum Protect Snapshot. This activity log records significant events such as completed commands and error messages. The IBM Spectrum Protect Snapshot activity log is distinct from the SQL Server error log. The *logfile* variable identifies the name to be used for the activity log generated by IBM Spectrum Protect Snapshot.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully qualified path; however, if you specify no path, the file is written to the directory where IBM Spectrum Protect Snapshot is installed.
- You cannot turn off IBM Spectrum Protect Snapshot logging activity. If you do not specify **/LOGFile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot to run operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays*|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option No can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify no, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of

the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OFFLOAD

Specify this option if, after the VSS snapshot is complete, you want to offload the transfer of the data from the IBM Spectrum Protect server to the system specified by the **REMOTEDSMAGENTNODE** parameter. This option is only valid when the **BACKUPDESTINATION** parameter is set to either TSM or BOTH. The default is to not offload data.

/Quiet

This parameter prevents status information from being displayed. This function does not affect the level of information that is written to the activity log.

/SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode that is used when you log on to the SQL Server. The integrated value specifies Windows authentication. The user ID you use to log on to Windows is the same ID you use to log on to the SQL Server. This option is the default value.

Use the *sqluserid* value to specify SQL Server user ID authorization. The user ID specified by the **/sqluserid** parameter is the ID you use to log on to the SQL Server. Any SQL user ID must have the SQL Server SYSADMIN fixed server role.

/SQLCHECKSum=No|Yes

The **/SQLCHECKSum** parameter is used to verify the integrity of a legacy database backup. Integrity checking is a process that validates the values in a file or configuration for unexpected changes. Values are verified between the current state and the baseline state.

You can specify the following values:

No

Do not enable integrity checking for a legacy database backup. This value is the default.

Yes

Enable integrity checking for a legacy database backup.

In the **Performance Properties** window of Microsoft Management Console, you can enable or disable the checksum option for all your legacy databases at once. You can override the global setting, and temporarily enable or disable the checksum option for a database backup, by setting this **SQLCHECKSum** parameter value to **Yes** or **No**.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL password that IBM Spectrum Protect Snapshot uses to log on to the SQL Server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user ID for this password must both be configured for SQL Server authentication.
- If you do not specify **/SQLPassword**, the default value is blank ("").
- If you specify **/SQLPassword**, but not *sqlpasswordname*, the default is also blank ("").
- This parameter is ignored if you use the **/SQLAUTH=INTEgrated** parameter with it.

/SQLSERVER=sqlserver-connection

The **/sqlserver** parameter specifies the SQL Server that Data Protection for SQL Server logs on to. The *sqlserver-connection* comprises the *sqlprotocol* and *sqlservername*. The *sqlprotocol* variable specifies the communication protocol to use and with this variable, you can specify an *sqlservername*. You can check the SQL connection by using the SQL Server Configuration Manager tool (under SQL Server Native Client Configuration client protocols). You can choose from the following protocols:

Table 39. SQL Server connection protocols		
Protocol Name	Description	Example Usage (with <i>sqlserver-connection</i> details)
lpc	Shared Memory	/sqlserver=lpc:<servername>\<instancename>
np	Named Pipes	/sqlserver=np:<servername>\pipe You can optionally specify a specific named pipe instance. For example, /sqlserver=np: \\hostname\pipe\pipe name By default, the pipe name is <i>sql\query</i> . If you connect to a named instance, the pipe name is typically in the following format: \\<servername>\pipe\MSSQL\$<instancename>\SQL\query
tcp	Transmission Control	/sqlserver=[tcp:]<servername>[\<instancename>][,port]
via	Virtual Interface Adapter	/sqlserver=via:<virtualservername>[\<instancename>]



Attention:

- For tcp protocols only, you have the option of defining a *port*. If you do not define a port, the default port value is the SQL default port 1433.
- For the via protocol, SQL Server supports this protocol only through SQL Server 2008 R2.
- To enable Data Protection for SQL Server to communicate with AlwaysOn Availability Group (AAG) instances, it is not possible to connect to the SQL Server using AAG listeners. For backup and restore operations, you must use the local SQL Server instance name (or instance name and port number) to communicate with the AAG. For AAG (or non-AAG instances), you can also specify non-default port numbers.

If you do not specify a protocol, Data Protection for SQL Server logs on to the SQL Server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL Server configurable option in the Data Protection for SQL Server configuration file. This is initially the local computer name.
- If you specify **/sqlserver** but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: . (local) That is, a period or the word *local* within parentheses.
- If the SQL Server is a member of a fail-over cluster, the CLUSTERNODE option in the IBM Spectrum Protect options file must have the value YES.
- If the SQL Server is not the default instance or is a member of a fail-over cluster, you must specify the name.
- The format of *sqlservername* depends on what type of instance it is and whether it is clustered or not:

Format	Instance?	Clustered?	Name required?
<i>local-computername</i>	default	no	no
<i>local-computername</i> \ <i>instancename</i>	named	no	yes
<i>virtualservername</i>	default	yes	yes
<i>virtualservername</i> \ <i>instancename</i>	named	yes	yes

localcomputername

The network computer name of the computer on which the SQL Server and Data Protection for SQL Server reside. The TCP/IP host name may not always be the same.

instancename

The name given to the named instance of the SQL Server that is specified during installation of the instance.

virtualservername

The name given to the clustered SQL Server that is specified during clustering service setup. This name is not the cluster or node name.

/SQLUser=sqlusername

The **/SQLUser** parameter specifies the name that IBM Spectrum Protect Snapshot uses to log on to the SQL Server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user ID for this password must both be configured for SQL Server authentication.
- The SQL user ID must have the SQL Server SYSADMIN fixed server role.
- If you do not specify **/SQLUser**, the default is *sa*.
- If you specify **/SQLUser** but not *sqlusername*, the default is also *sa*.
- This parameter is ignored if you use the **/SQLAUTH=integrated** parameter with it.

/USEALWAYSONnode

Specify this parameter to back up standard databases on SQL Server 2012 and later versions by using the AlwaysOn node. By setting this parameter, you can back up all availability databases and standard databases under a single node to help you to manage your database backups more easily. By default, SQL Server 2012 and later version availability databases are backed up to the AlwaysOn node.

Backup examples

The following examples are provided to show how the **backup** command can be entered with various parameters and options.

If you want to use the **backup** command from the command-line interface, the following examples are provided to help model the command syntax:

- To complete a full backup of a database, enter the following command:

```
tdpsqlc backup
```

- To complete a full backup of all standard databases, enter the following command:

```
tdpsqlc backup * full /EXCLUDEALwaysondbs
```

- To complete a log backup of all availability databases, enter the following command:

```
tdpsqlc backup * log /EXCLUDESTandarddbbs
```

- For a more complex example, consider the following scenario: There are three AlwaysOn Availability Groups. The first availability group is called *AG01* with the following databases:

- AlwaysOn Availability Database called *AlwaysOnLegacyDB1*
- AlwaysOn Availability Database called *AlwaysOnLegacyDB3*

The second availability group is called *AG03* with the following AlwaysOn Availability Database: *AlwaysOnLegacyDB2*. The third availability group is called *AG04* with the following databases:

- AlwaysOn Availability Database called *AlwaysOnLegacyDB5*
- AlwaysOn Availability Database called *AlwaysOnLegacyDB6*
- Standard database that is called *SQL_DB1*
- Standard database that is called *SQL_DB2*

To complete a full backup with a database list that matches both standard and availability databases, but excluding standard databases, enter the following command:

```
C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup AlwaysOnLegacy*,SQL*
full /backupdest=TSM /backupmeth=legacy /EXCLUDEStandarddbs
```

- When you use the **/AAGName** parameter to filter the databases that are backed up, refer to the following scenario with the examples: There are two AlwaysOn Availability Groups. The first availability group is called *AG01* with the following databases:
 - AlwaysOn Availability Database called *AlwaysOnLegacyDB1*
 - AlwaysOn Availability Database called *AlwaysOnLegacyDB3*

The second availability group is called *AG04* with the following databases:

- AlwaysOn Availability Database called *AlwaysOnLegacyDB5*
- AlwaysOn Availability Database called *AlwaysOnLegacyDB6*

When you enter a **backup** command for all databases, but use the **/AAGName** parameter to include only databases from *AG01* in the backup, enter the following command:

```
C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup * full /backupdest=TSM
/backupmeth=legacy /AAGName=AG01
```

When you enter a **backup** command for a database list with wildcards, but use the **/AAGName** parameter to include only databases from *AG04* in the backup, enter the following command:

```
C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup AlwaysOn*,SQL* full
/backupdest=TSM /backupmeth=legacy /AAGName=AG04
```

When you enter a **backup** command for a database list with wildcards, but do not match all databases from the specified AlwaysOn Availability Group, enter the following command:

```
C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc back *DB5 full /backupdest=TSM
/backupmeth=legacy /AAGName=AG04
```

- To complete a differential backup with a database list that matches both standard and availability databases, but excluding availability databases, enter the following command:

```
C:\Program Files\tivoli\tsm\TDPSql>tdpsqlc backup AlwaysOnLegacy*,SQL*
diff /EXCLUDEAlwaysondbs
```

Delete backup command

Use the **delete backup** command to delete a VSS backup of a SQL Server database.

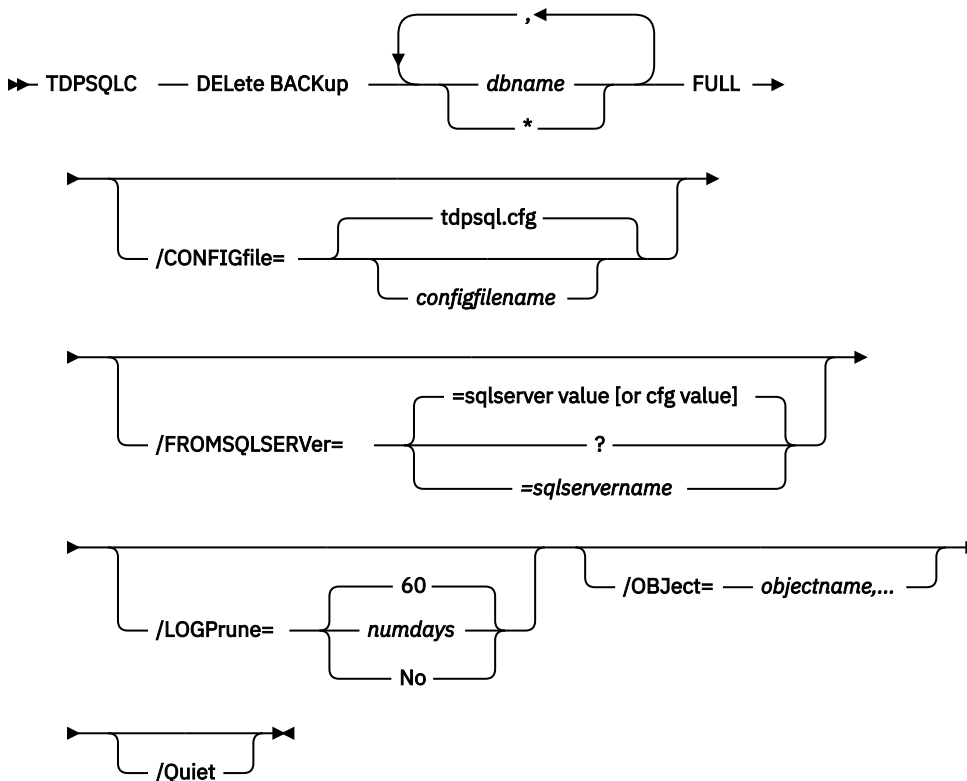
You must have local registry rights (for all versions of SQL Server) to run an IBM Spectrum Protect Snapshot for SQL Server delete backup.

- If you delete multiple LOCAL snapshots that are stored on SAN Volume Controller or Storwize family Space Efficient volumes (SEV), you must do so in the same order in which you created the snapshots. That is, you must delete the oldest one first, followed by the second oldest. Failure to delete them in this order can cause removal of other snapshots of the same source.
- If you mount a local VSS COPY type backup as a snap of a snap, the snap of a snap volume is also deleted along with the VSS backup.

Delete Backup syntax

Use the **delete backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Delete Backup positional parameters

Positional parameters immediately follow the **delete backup** command and precede the optional parameters.

The following positional parameters specify the backup to delete:

* | *dbname*

*

Delete the active backups of all databases.

dbname

Delete a backup of the specified database. The active backup is deleted unless you specify a different backup with the **/object** optional parameter.

Multiple entries are separated by commas. If separated by commas, make sure that there is no space between the comma and the database name. If any database name contains blanks, enclose the database name in double quotation marks.

The following positional parameter specifies the type of delete backup to run:

FULL

Delete full database backups.

COPYFULL

Delete copy-only full database backups.

Delete Backup optional parameters

Optional parameters follow the **delete backup** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for SQL Server configuration file that contains the values to use for a **delete backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for SQL Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpsql.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

See “Set positional parameters” on page 250 for descriptions of available configuration parameters.

/FROMSQLSERVER=*server-name*

Use the **/fromsqlserver** parameter to specify the name of the SQL Server where the original backup was done. This parameter is necessary only when the name of the SQL Server to delete from, as determined by the **/sqlserver** parameter, is different from the name of the SQL Server that the backup objects were created from. The default value is the **/sqlserver** value or the value that is set in the IBM Spectrum Protect Snapshot configuration file.

Considerations:

- If the two SQL Server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/LOGFile=*logfile*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for SQL Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for SQL Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpsql.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for SQL Server to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=*numdays***|No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJECT=objectname,...

Use the **/object** parameter to specify the names of backup objects you want to delete. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot for SQL Server.

Use the IBM Spectrum Protect Snapshot for SQL Server query `fcmsql * /all` command to view the names of all available backup objects. This parameter specifies that only particular backup objects for the specified SQL databases and backup object type is to be deleted. The *objectname* variable specifies the names of the backup objects you want to delete. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot for SQL Server.

/QUERYNode=DP | ALWAYSON

Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard Data Protection for SQL node or the AlwaysOn node. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/Quiet

This parameter prevents status information from being displayed. This function does not affect the level of information that is written to the activity log.

Delete backup example

This output example provides a sample of the text, messages, and process status that displays when you use the **delete backup** command.

In this example, the `tdpsqlc delete backup xivdb1 full` command deletes a full backup of database *xivdb1*. The following output is displayed:

```
Connecting to SQL Server, please wait...
Querying for Backups ....
Backup(s) to be deleted:
  xivdb1 : VSS : full : 02/10/2014 10:03:29
VSS Delete backup operation completed with rc = 0
  Files Examined      : 1
  Files Completed     : 1
  Files Failed        : 0
  Total Bytes         : 0
```

Help command


Use the **tdpsqlc help** command to display help for IBM Spectrum Protect Snapshot for SQL Server commands.

This command lists one or more commands and their parameters. When you use a language that is not English, you might be required to set the width of your screen display to a value greater than 80 characters. The wider setting displays the entire help description in one screen. For example, set the screen width to 100 characters.

Help syntax

Use the **help** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Help positional parameters

Positional parameters immediately follow the **help** command. There are no optional parameters with this command.

Use the help command to display the syntax of all or selected IBM Spectrum Protect Snapshot commands by using a textual notation.

Help uses the following notation:

[a]

a is optional; *a* might occur zero or one time

{a | b}

Select either *a* or *b*, but not both

{a}+

a must occur at least one time

{a}*

a might occur zero or more times

(a)

Comments that are not part of the command

UPPERCASE

Minimum abbreviation (which you can also enter in lowercase)

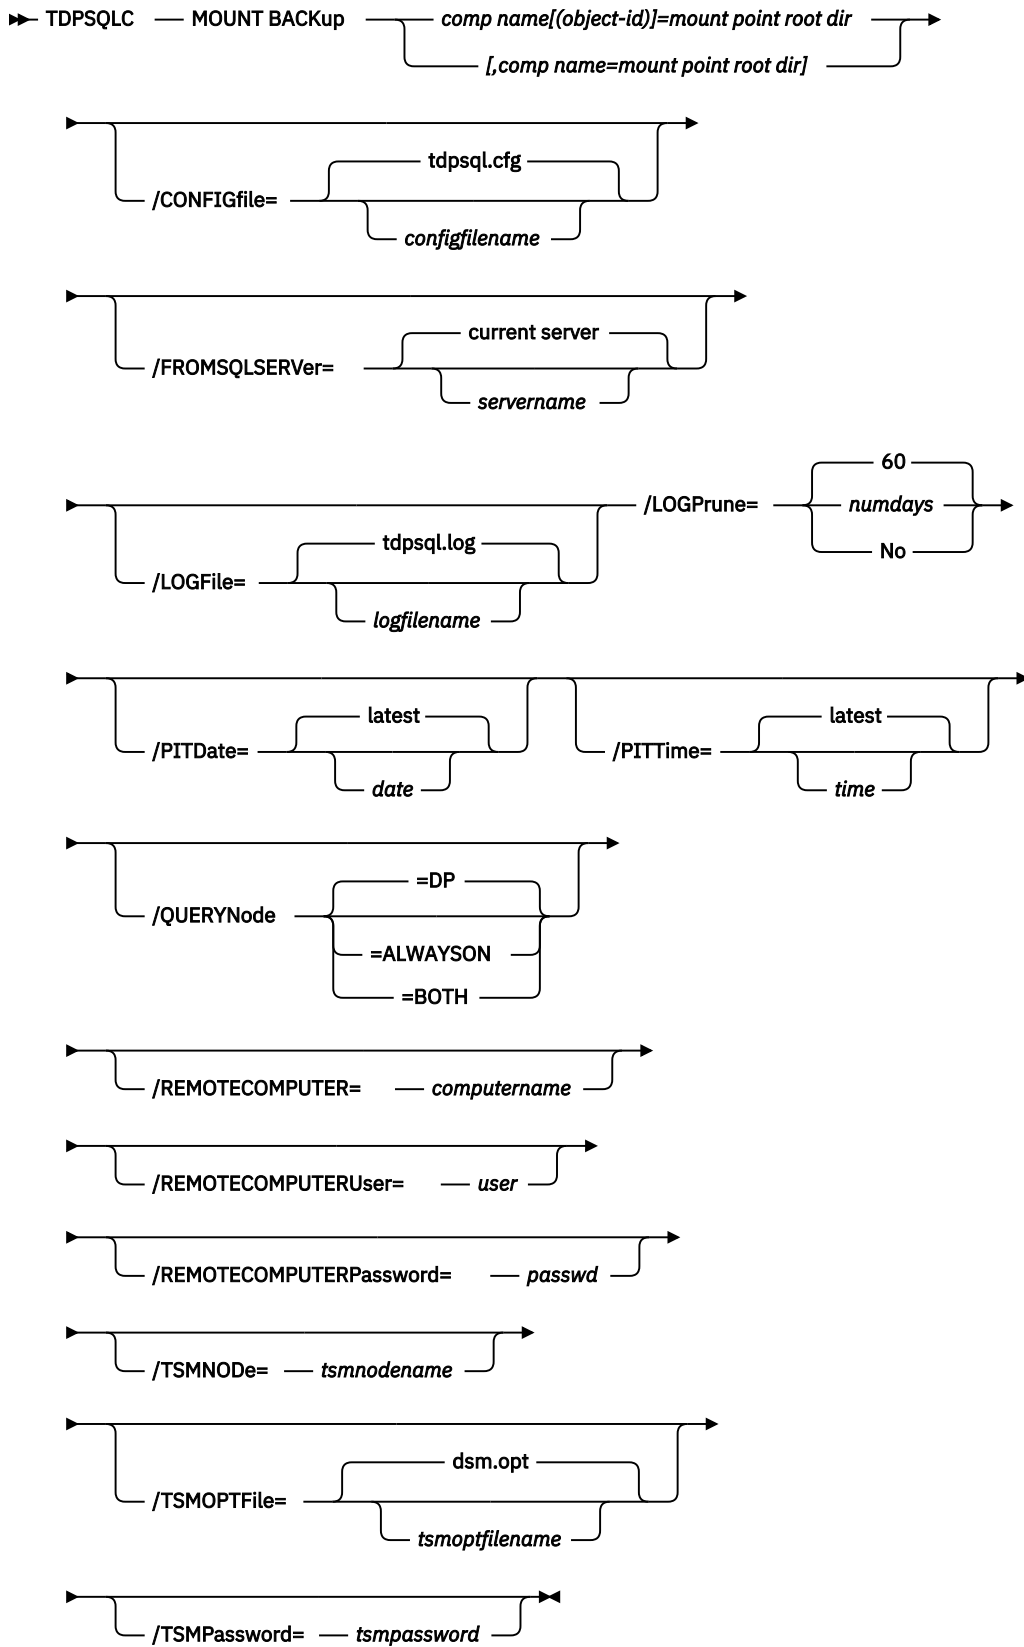
Mount Backup command

Use the **mount backup** command to mount backups that are managed by IBM Spectrum Protect Snapshot for SQL Server.

Mount Backup syntax

Use the **mount backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Mount Backup positional parameter

The positional parameters immediately follow the **mount backup** command and precede the optional parameters.

The following positional parameters specify the objects to mount:

component name[(object-id)]=mount point root dir[,component name=mount point root dir]

component name[(object-id)]

Specify the backup of a local SQL Server database or storage group.

mount point root dir

Specify the absolute path to the directory where the snapshots are going to be displayed as mount point directories. The directory must be empty. If not empty, an error is reported.

The list must contain all non-qualified objects or all qualified objects. The list cannot contain a combination of non-qualified objects and qualified objects. Specify the list by using the following syntax:

```
mount backup object-1[(object-1-id)]= mount-point-1[,object-2[(object-2-id)]
=mount-point-2...]
```

Mount Backup optional parameters

Optional parameters follow the **mount backup** command and positional parameters.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for SQL Server configuration file that contains the values to use for a **mount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for SQL Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpsql.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpsql.cfg"
```

/FROMSQLSERVER=server-name

Use the **/fromsqlserver** parameter to specify the name of the server where the original backup was done. The default is the local server.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for SQL Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpsql.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain

number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/PITDate=date

Use the **/pitdate** parameter with the **/pittime** parameter to establish a point in time for which you want to mount the latest version of your backups. Backups that were backed up on or before the date and time you specified, and that were not deleted before the date and time you specified, are processed. Backup versions that you create after this date and time are ignored. Specify the appropriate date in the *date* variable; use the same format that you selected with the **DATEFORMAT** option in the IBM Spectrum Protect Snapshot for SQL Server options file.

If the *date* or the *time* is not specified, then no date and time are established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup that is selected after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the mount date and time as the current date at the specified *time*.

/PITTime=time

Use the **/pittime** parameter with the **/pitdate** option to establish a point in time for which you want to mount the latest version of your backups. Files or images that were backed up on or before the date and time you specify, and that were not deleted before the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify the **/pitdate** parameter. Specify the appropriate time in the *time* variable; use the same format that you selected with the **TIMEFORMAT** option in the IBM Spectrum Protect Snapshot for SQL Server options file.

If the *date* or the *time* is not specified, then no date and time are established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup that is selected after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the mount date and time as the current date at the specified *time*.

/QUERYNode=DP | ALWAYSON | BOTH

Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard Data Protection for SQL Server node, the AlwaysOn node, or both nodes. To mount a backup that is using the AlwaysOn node (for AlwaysOn Availability databases), specify `/QUERYNode = ALWAYSON`.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to mount the data.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter `Domain\User`. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=password

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnode name

Use the *tsmnode name* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (`dsm.opt`). This parameter overrides the value in the IBM Spectrum Protect options file if `PASSWORDACCESS` is set to `PROMPT`. This parameter is not valid when `PASSWORDACCESS` is set to `GENERATE` in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire `/tsmoptfile` parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\dsm.opt"
```

The default is `dsm.opt`.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified `PASSWORDACCESS GENERATE` in the IBM Spectrum Protect Snapshot options file (`dsm.opt`), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time that IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when `PASSWORDACCESS GENERATE` is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If `PASSWORDACCESS PROMPT` is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Mount Backup examples

These output examples provide a sample of the text, messages, and process status that displays when you use the `mount backup` command.

Examples:

```
TDPSTQLC MOUNT BACKup SQL-DB-1=K:\MP-dir
TDPSTQLC MOUNT BACKup SQL-DB-1(20120523070512)=L:\MP-dir
```

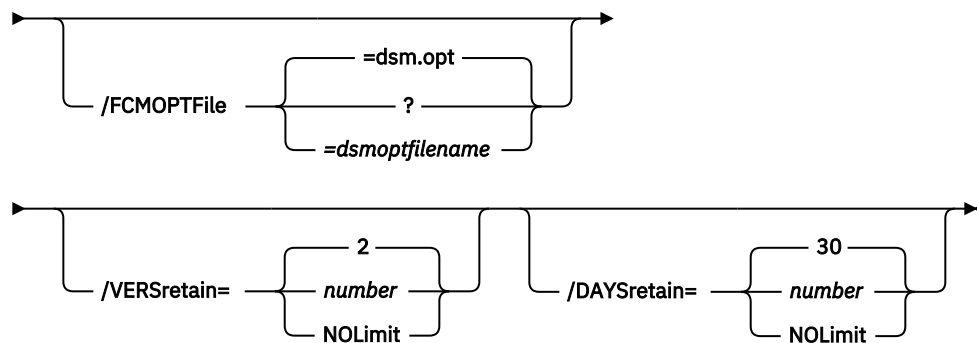
Policy commands for IBM Spectrum Protect Snapshot for SQL

Create Policy

This command is used to create a policy.

TDPSTQLC command: CREATE POLIcy

➤ TDPSTQLC — CREATE POLIcy — *policy_name* →



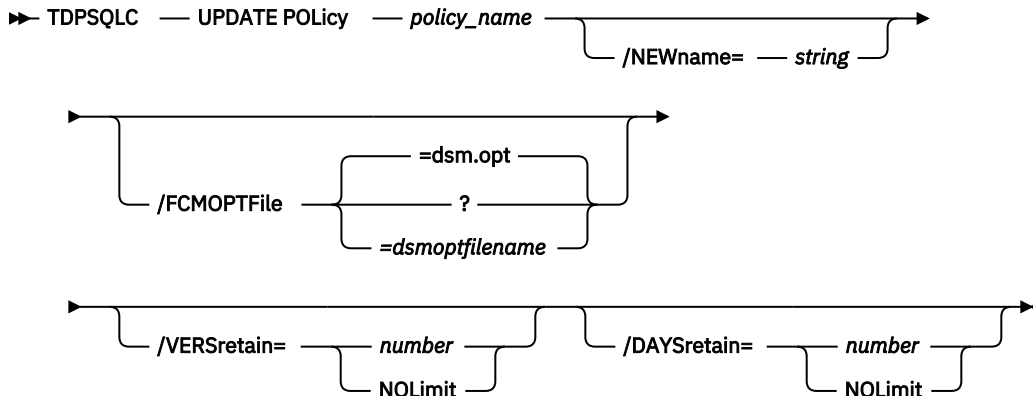
Parameters:

- *policy_name* (required): Specifies the name of the policy that is being created. To create a policy, the policy name must be unique.
- **VERsretain**: Specifies the number of snapshot versions to retain (1 - 9999). You can also specify **NOLimit** to represent an unlimited number of snapshot versions to retain.
- **DAYsretain**: Specifies the number of days to retain a snapshot (0 - 9999). You can also specify **NOLimit** to represent an unlimited number of days to retain snapshot versions.

Update Policy

This command is used to update or modify the retention parameters of an existing policy.

TDPSQLC command: UPDATE POLIcy



Parameters:


- **NEWname**: Specifies the new name of the policy, if the name is being updated. The policy name must be unique.
- *policy_name* (required): Specifies the name of the policy that is being updated.
- **VERSretain**: Specifies the number of snapshot versions to retain (1 - 9999). You can also specify **NOLimit** to represent an unlimited number of snapshot versions to retain.
- **DAYSretain**: Specifies the number of days to retain a snapshot (0 - 9999). You can also specify **NOLimit** to represent an unlimited number of days to retain snapshot versions.

Copy Policy

This command is used to copy an existing policy to a new policy.

TDPSQLC command: COPY POLIcy

```
TDPSQLC — COPY POLIcy — existing_policy_name — new_policy_name —
```



Parameters:

- *existing_policy_name* (required): Specifies the name of the policy that is being copied.
- *new_policy_name* (required): Specifies the name of the new policy. The policy name must be unique.

Query Policy

This command is used to list the attributes of a policy.

TDPSQLC command: Query POLIcy

```
TDPSQLC — Query POLIcy — * —
```

Parameters: * (required) Specifies all policies are to be queried. The results of the query are displayed as follows:


Connecting to SQL Server, please wait...

Policy	Number of snapshots to keep	Days to keep a snapshot
FCMPOL	3	60
STANDARD	2	30

Delete Policy

This command is used to delete a policy.

TDPSQLC command: DElete POLicy



Parameter:

- *policy_name* (required): Specifies the name of the policy that is being deleted.

Query FCM command

Use the **query fcm** command to display IBM Spectrum Protect Snapshot information.

This command displays the following information:

- Compression mode
- Active policy set
- Default management class

This command can also display a list of backups that match the databases that are entered.

Active and inactive objects can be displayed. However, only the active backup objects are displayed by default. To include inactive backup versions in the list, use the **/all** optional parameter.

Query FCM example

Use the **query fcm** command to return output about the server and other information:

```
IBM Spectrum Protect server Connection Information
-----
Nodename ..... MALTA_EXC
Network Host Name of Server ..... FVTSERIES10
TSM API Version ..... Version 7, Release 1, Level 3.0

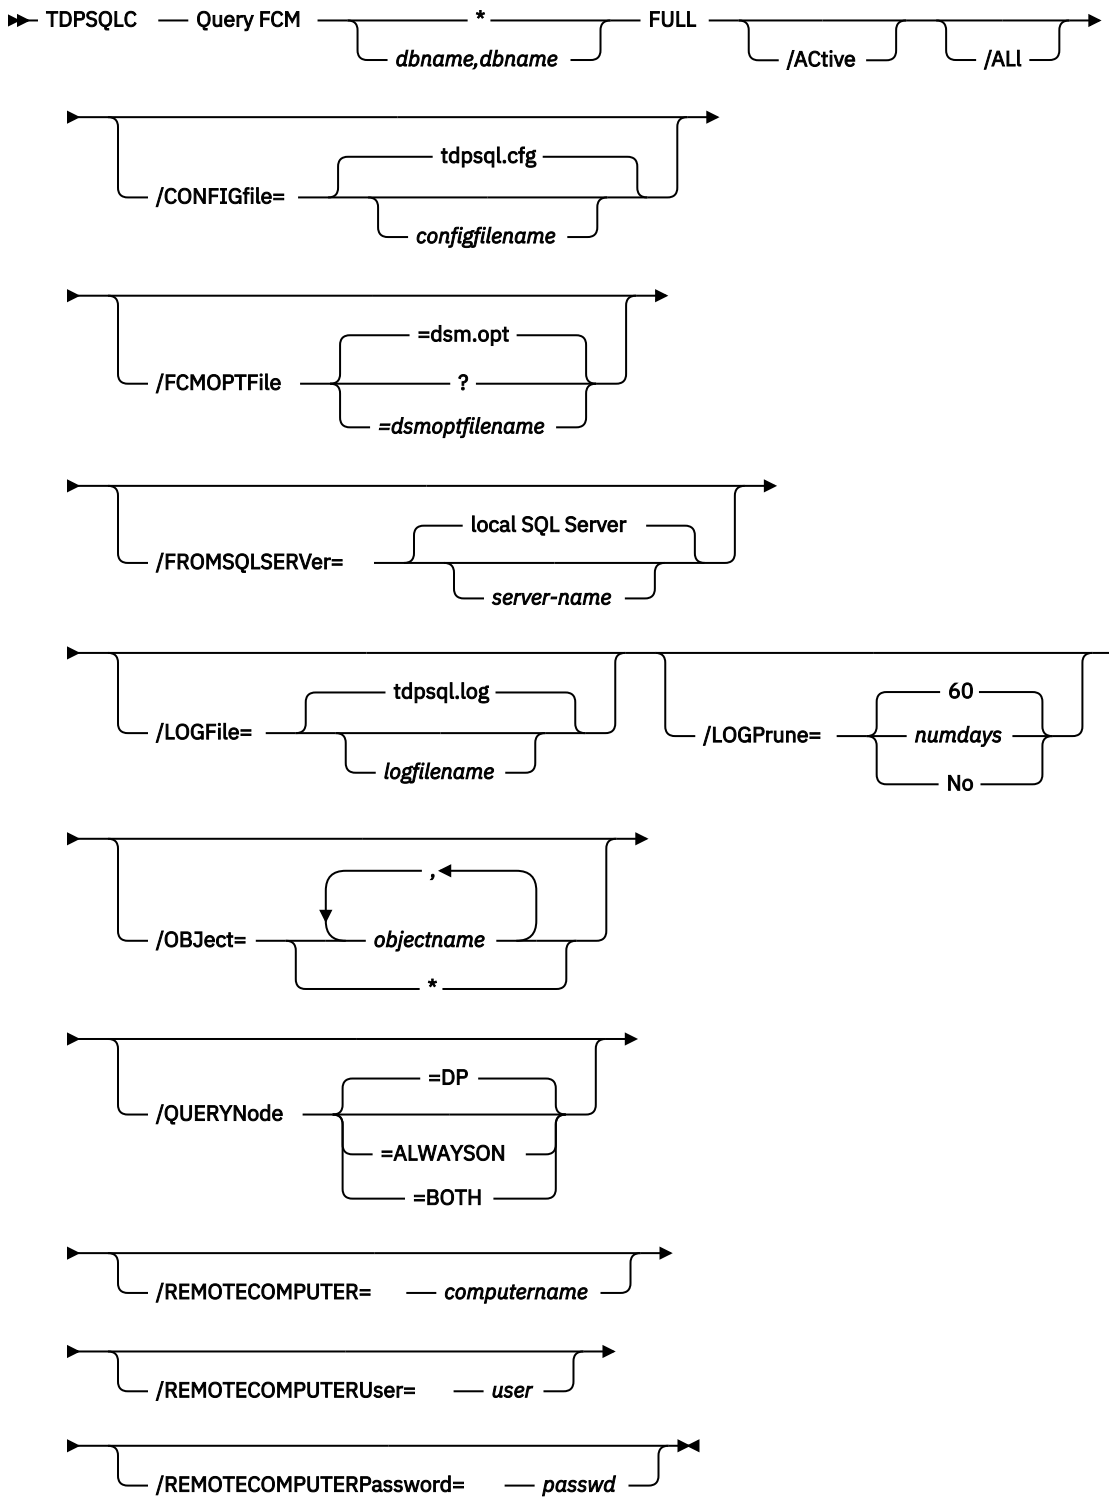
Server Name ..... FVTSERIES10_SERVER1_622GA
Server Type ..... Windows
Server Version ..... Version 7, Release 1, Level 3.0
Compression Mode ..... Client Determined
Domain Name ..... FCM_PDEXC
Active Policy Set ..... STANDARD
Default Management Class ..... STANDARD

Completed
```


Query FCM syntax

Use the **query FCM** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Query FCM positional parameters

Positional parameters immediately follow the **query FCM** command and precede the optional parameters.

The following positional parameters specify the object to query. If none of these positional parameters are specified, only the IBM Spectrum Protect Snapshot API and IBM Spectrum Protect Snapshot information is displayed:

* | **dbname**

dbname1, ..., dbnameN

Query all backup objects for the specified database. Multiple entries are separated by commas. where *dbname* can be a database name.

The following positional parameters specify the type of backup to query. If this parameter is not specified, all backup types are displayed:

FULL

Query only full backup types.

COPY

Query only copy backup types.

INCR

Query only incremental backup types.

DIFF

Query only differential backup types.

Query FCM optional parameters

Optional parameters follow the **query FCM** command and positional parameters.

/Active

Use the **/active** parameter to display active backup objects only. This parameter is the default.

/All

Use the **/all** parameter to display both active and inactive backup objects. If the **/all** parameter is not specified, only active backup objects are displayed.

/CONFIGfile=configfilename

The **/configfile** parameter specifies the name of the IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file, which contains the values for the IBM Spectrum Protect Snapshot configurable options. See [“Set command” on page 306](#) for details on the content of the file.

Considerations:

- The *configfilename* variable can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where IBM Spectrum Protect Snapshot for Microsoft SQL Server is installed.
- If *configfilename* includes spaces, enclose it in double quotation marks.
- If you do not specify **/configfile**, the default value is `tdpsql.cfg`.
- If you specify **/configfile** but not *configfilename*, the default value `tdpsql.cfg` is used.

/FCMOPTfile=dsmoptfilename

The **/fcmoptfile** parameter specifies the IBM Spectrum Protect Snapshot options file to use.

Considerations:

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/fcmoptfile**, the default value is `dsm.opt`.
- If you specify **/fcmoptfile** but not *dsmoptfilename*, the default is also `dsm.opt`.

`/FROMSQLSERVER=sqlservername`

For **query FCM**, the **`/fromsqlserver`** parameter specifies the SQL Server that backup objects were backed up from. This parameter is necessary only when the name of the SQL Server to query, as determined by the **`/sqlserver`** parameter, is different from the name of the SQL Server that the backup objects were created from. The default value is the **`/sqlserver`** value or the value that is set in the IBM Spectrum Protect Snapshot for Microsoft SQL Server configuration file.

Considerations:

- If the two SQL Server names are different, you must use this parameter even if **`/fromsqlserver`** was a non-clustered default instance.
- After you restore a SQL database to a different SQL Server, the logins of the SQL database might not match the logins for the different SQL Server. If appropriate, you can use the SQL stored procedure `SP_CHANGE_USERS_LOGIN` to find and correct such SQL login mismatches.

`/LOGFile=logfilename`

Use the **`/logfile`** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for Microsoft SQL Server.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for Microsoft SQL Server installation directory.

If the *logfilename* variable includes spaces, enclose the entire **`/logfile`** parameter entry in double quotation marks For example:

```
/LOGFile="c:\Program Files\mytdpsql.log"
```

You cannot turn IBM Spectrum Protect Snapshot for Microsoft SQL Server logging activity off. If you do not specify **`/logfile`**, log records are written to the default log file. The default log file is `tdpsql.log`.

Attention: When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for Microsoft SQL Server to run operations, use the **`/logfile`** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

`/LOGPrune=numdays|No`

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **`/LOGPrune`** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **`/LOGPrune`**, the default value is that specified by the **`logprune`** configurable option in the configuration file. The default value is 60.
- If you specify **`/LOGPrune`**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.

- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJECT=* | objectname,...

For **restore** and **inactivate** operations, **/object** specifies that only particular backup objects for the specified SQL databases and backup object type (if specified) be restored or deactivated. For **query** operations, **/object** includes particular objects and object types in the display. The *objectname* variable specifies the names of the backup objects you want to restore or deactivate. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot. Use **query** to view the names of backup objects. You can specify the asterisk (*) wildcard character in *objectname* to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/QUERYNode=DP | ALWAYSON | BOTH

Specify whether you want to query standard databases from SQL Server 2012 and later versions that are backed up from a standard IBM Spectrum Protect Snapshot for Microsoft SQL Server node, the AlwaysOn node, or both nodes. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to query the data that is backed up.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

Query Managedcapacity command

Use the **Query Managedcapacity** command to assist with storage planning by determining the amount of managed capacity in use.

Purpose

The **query managedcapacity** command displays capacity that is related information about the volumes that are represented in local inventory that is managed by IBM Spectrum Protect Snapshot. This command is valid for all Windows operating systems that are supported by IBM Spectrum Protect Snapshot.

The capacity that is displayed includes deactivated backups, that is, backups which have not expired, on the IBM Spectrum Protect server. Once a deleted backup has been expired by the IBM Spectrum Protect server, the capacity that is displayed no longer contains capacity for the deleted backup.

TDPSQLC command

➤ TDPSQLC — Query MANAGEDCAPacity 

Parameters

/DETAILED

Results in a detailed listing of snapped volumes. If this option is not specified, then only the total capacity is displayed.

SQL Server 2008 example

Query the total managed capacity of SQL Server 2008 data that is represented in the local inventory with a detailed listing of snapped volumes.

Command: `tdpsqlc query managedcapacity /detailed`

```
Total Managed Capacity : 63.99 GB (68,706,877,440 bytes)
```

```
Volume      : H:  
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

```
Volume      : I:  
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

```
Volume      : Q:  
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

```
Volume      : N:  
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

Query SQL command

The **query sql** command queries the local SQL Server to return general information and status about the SQL Server, databases, and VSS components.

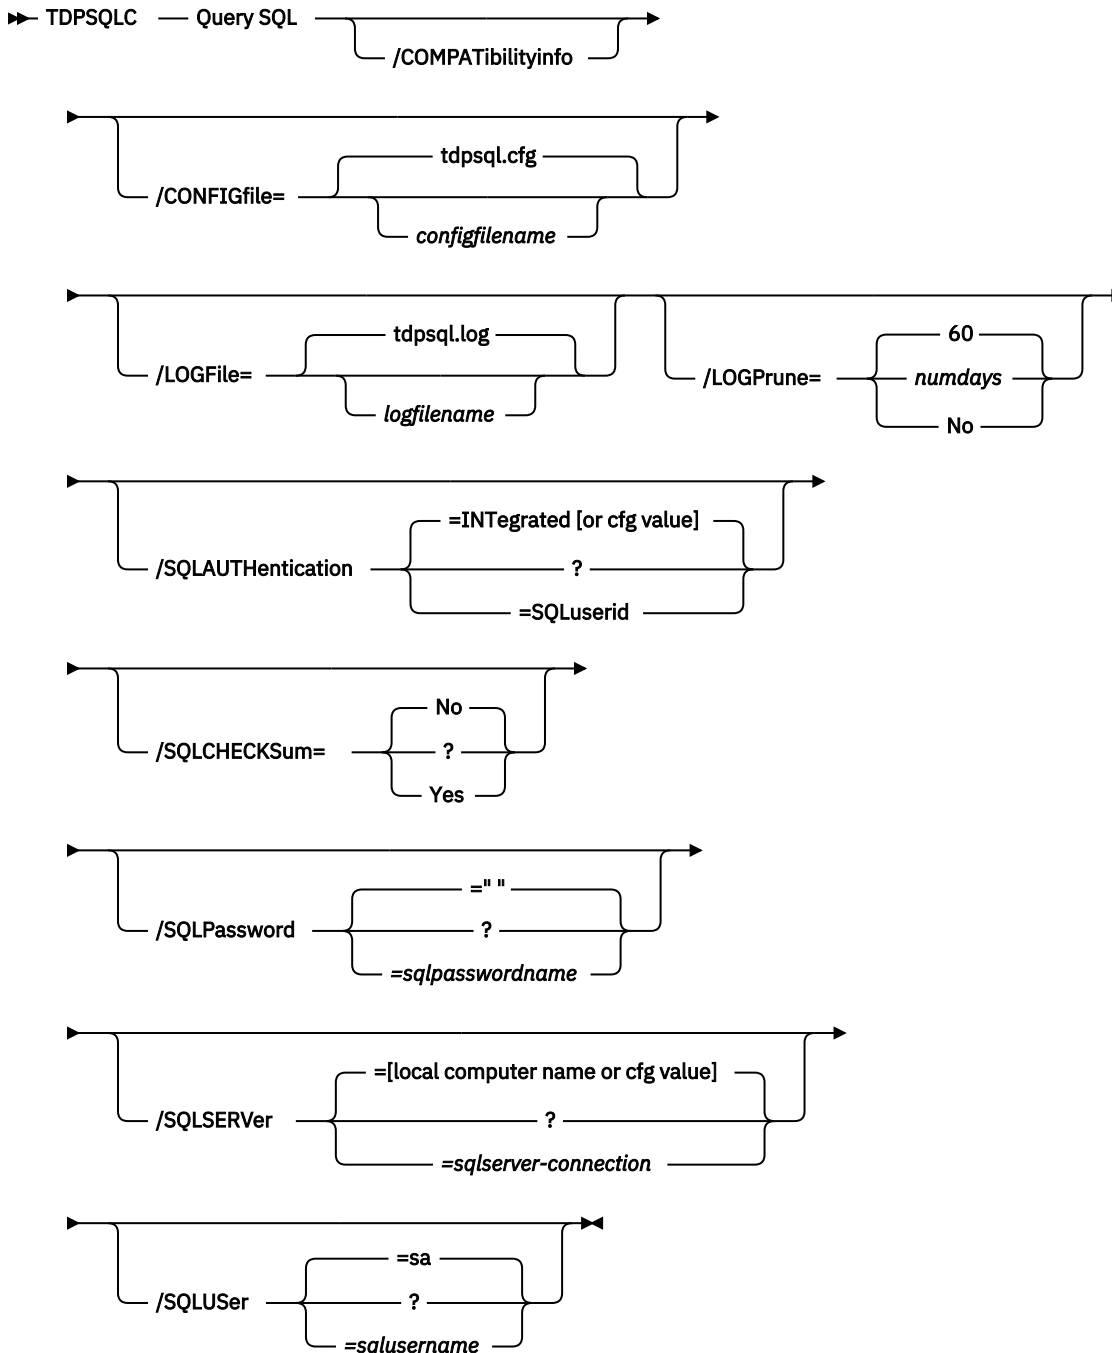
Use the **query sql** command to return the following information:

- SQL Server information:
 - SQL Server name and version
 - Database name
 - Database data space allocated
 - Database space that is used
 - Database log space allocated
 - Database log space used
 - Database options that are set (SELECT INTO / BULK COPY, TRUNCATE LOG ON CHECKPOINT, and other options.)
- VSS information:
 - Writer Name
 - Local DSMAgent Node
 - Remote DSMAgent Node
 - Writer Status (online, offline)
 - Number of selectable components
- If you specify **/compatibilityinfo**:
 - Server clustering state
 - Database compatibility level

Query SQL syntax

Use the **query sql** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Query SQL positional parameters

Positional parameters immediately follow the **query** command and precede the optional parameters.

Specify one of the following when you issue an IBM Spectrum Protect Snapshot for SQL Server **query** command:

Query SQL * | dbname,...

This displays information about the current SQL Server. The *dbname* variable specifies databases on the current SQL Server to display information about.

Query SQL optional parameters

Optional parameters follow the **query sql** command and positional parameters.

/COMPATibilityinfo

For **query** operations, this parameter displays information that is related to the compatibility of a backup object with a SQL Server. Certain SQL Server configuration options must be compatible before you can restore a backup object to a SQL Server. When you specify this parameter, SQL and IBM Spectrum Protect Snapshot for SQL Server configuration information is listed to help you determine whether a backup object is correct for a SQL Server.

Considerations:

- Compatible generally means identical. However, if you use a binary sort order for both the SQL Server and the backup object, the code pages might be different, although the interpretation of individual character values might result in different characters that are displayed or printed.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for SQL Server configuration file that contains the values to use for a **query sql** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for SQL Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpsql.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

See [Set positional parameters](#) for descriptions of available configuration parameters.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for SQL Server. The *logfilename* variable identifies the name of the activity log file. If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for SQL Server installation directory. If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpsql.log`. The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for SQL Server to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode that is used when you log on to the SQL Server. The **integrated** value specifies Windows authentication. The user ID you use to log on to Windows is the same ID you use to log on to the SQL Server. This option is the default value. Use the *sqluserid* value to specify SQL Server user ID authorization. The user ID specified by the **/sqluserid** parameter is the ID you use to log on to the SQL Server. Any SQL user ID must have the SQL Server SYSADMIN fixed server role.

/SQLCHECKSum=No|Yes

Use the **SQLCHECKSum** parameter to verify the integrity of a legacy database backup.

You can specify the following values:

No

Do not enable the checksum option for a legacy database backup. This option is the default option.

Yes

Enable the checksum option to verify that a legacy database backup is consistent and correct.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL password that IBM Spectrum Protect Snapshot uses to log on to the SQL Server that objects are backed up from or restored to.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user ID for this password must both be configured for SQL Server authentication.
- If you do not specify **/sqlpassword**, the default value is blank (" ").
- If you specify **/sqlpassword** but not *sqlpasswordname*, the default is also blank (" ").
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVER=sqlserver-connection

The **/sqlserver** parameter specifies the SQL Server that Data Protection for SQL Server logs on to. The *sqlserver-connection* comprises the *sqlprotocol* and *sqlservername*. The *sqlprotocol* variable specifies the communication protocol to use and with this variable, you can specify an *sqlservername*. You can check the SQL connection by using the SQL Server Configuration Manager tool (under SQL Server Native Client Configuration client protocols). You can choose from the following protocols:

<i>Table 40. SQL Server connection protocols</i>		
Protocol Name	Description	Example Usage (with <i>sqlserver-connection</i> details)
lpc	Shared Memory	/sqlserver=lpc:<servername>\<instancename>

Table 40. SQL Server connection protocols (continued)		
Protocol Name	Description	Example Usage (with <i>sqlserver-connection</i> details)
np	Named Pipes	<p>/sqlserver=np:<servername>\pipe</p> <p>You can optionally specify a specific named pipe instance. For example, /sqlserver=np: \\hostname\pipe\pipe name</p> <p>By default, the pipe name is <i>sql\query</i>. If you connect to a named instance, the pipe name is typically in the following format: \\<servername>\pipe\MSSQL\$<instancename>\SQL\query</p>
tcp	Transmission Control	/sqlserver=[tcp:]<servername>[\<instancename>][,port]
via	Virtual Interface Adapter	/sqlserver=via:<virtualservername>[\<instancename>]



Attention:

- For tcp protocols only, you have the option of defining a *port*. If you do not define a port, the default port value is the SQL default port 1433.
- For the via protocol, SQL Server supports this protocol only through SQL Server 2008 R2.
- To enable Data Protection for SQL Server to communicate with AlwaysOn Availability Group (AAG) instances, it is not possible to connect to the SQL Server using AAG listeners. For backup and restore operations, you must use the local SQL Server instance name (or instance name and port number) to communicate with the AAG. For AAG (or non-AAG instances), you can also specify non-default port numbers.

If you do not specify a protocol, Data Protection for SQL Server logs on to the SQL Server according to the first protocol that becomes available.

Considerations:

- The default value is the value specified by the SQL Server configurable option in the Data Protection for SQL Server configuration file. This is initially the local computer name.
- If you specify **/sqlserver** but not *sqlservername*, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: . (local) That is, a period or the word *local* within parentheses.
- If the SQL Server is a member of a fail-over cluster, the CLUSTERNODE option in the IBM Spectrum Protect options file must have the value YES.
- If the SQL Server is not the default instance or is a member of a fail-over cluster, you must specify the name.
- The format of *sqlservername* depends on what type of instance it is and whether it is clustered or not:

Format	Instance?	Clustered?	Name required?
<i>local-computername</i>	default	no	no
<i>local-computername</i> <i>\instancename</i>	named	no	yes
<i>virtualservername</i>	default	yes	yes
<i>virtualservername</i> <i>\instancename</i>	named	yes	yes

localcomputername

The network computer name of the computer on which the SQL Server and Data Protection for SQL Server reside. The TCP/IP host name may not always be the same.

instancename

The name given to the named instance of the SQL Server that is specified during installation of the instance.

virtualservername

The name given to the clustered SQL Server that is specified during clustering service setup. This name is not the cluster or node name.

/SQLUser=sqlusername

The **/sqluser** parameter specifies the name that IBM Spectrum Protect Snapshot uses to log on to the SQL Server.

Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user ID for this password must both be configured for SQL Server authentication.
- The SQL user ID must have the SQL Server SYSADMIN fixed server role.
- If you do not specify **/sqluser**, the default is sa.
- If you specify **/sqluser** but not *sqlusername*, the default is also sa.
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

Query SQL example

This output example provides a sample of the text, messages, and process status that displays when you use the **query SQL** command.

In this example, the **tdpsqlc query sql** command queried the local SQL Server to return general information and status about the SQL Server, databases, and VSS components. The following output is displayed:

```

Connecting to SQL Server, please wait...

SQL Server Information
-----

SQL Server Name      ..... VADER
SQL Server Version   ..... 10.0.1600 (SQL Server 2008)

Volume Shadow Copy Service (VSS) Information
-----

Writer Name          : SqlServerWriter
Local DSMAgent Node  : VADER
Remote DSMAgent Node :
Writer Status        : Online
Selectable Components : 13

```

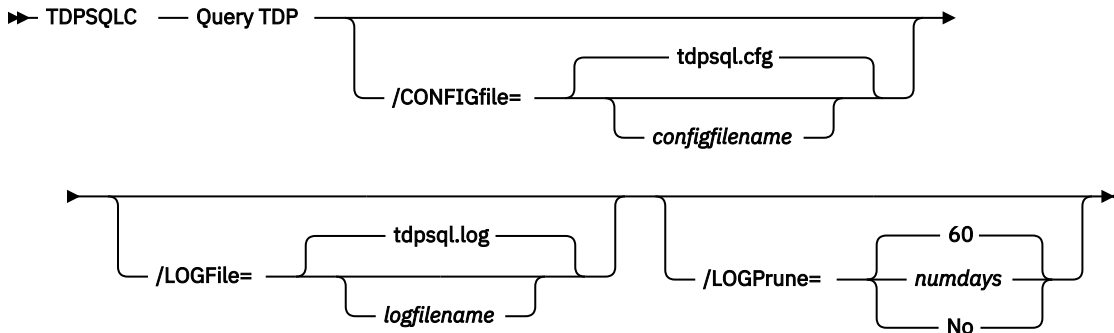
Query TDP command

Use the **query tdp** command to query a list of the current values set in the configuration file for IBM Spectrum Protect Snapshot for SQL Server.

Query TDP syntax

Use the **query TDP** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Query TDP optional parameters

Optional parameters follow the **query TDP** command.

/CONFIGfile=logfile

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot for SQL Server configuration file that contains the values to use for a **query tdp** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot for SQL Server installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpsql.cfg`

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

See [“Set positional parameters”](#) on page 307 for descriptions of available configuration parameters.

/LOGFile=logfile

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot for SQL Server.

The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot for SQL Server installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpsql.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for SQL Server to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This

function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option No can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify no, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

Query TDP example

This output example provides a sample of the text, messages, and process status that displays when you use the **query TDP** command.

In this example, the **tdpsqlc query tdp** command queried a list of the current values that are set in the configuration file for IBM Spectrum Protect Snapshot. The following output is displayed:

```
IBM Tivoli Storage FlashCopy Manager configuration settings
-----
CONFIGfile..... tdpsql.cfg
LOGFile ..... tdpsql.log
LOGPrune ..... 60
```

Restore command

Use this command to restore one (or more) SQL databases from storage that is managed by IBM Spectrum Protect Snapshot to a SQL Server.

Considerations:

- You cannot restore SQL databases currently in use. By placing SQL databases to be restored in single-user mode, you can avoid attempting such restores. If you are restoring the master database, you must start the SQL Server in single-user mode by using the -m SQL SERVER startup option.

Note:

1. The single user of the SQL databases or server must be the same user that IBM Spectrum Protect Snapshot uses to log on to the SQL Server for the restore.

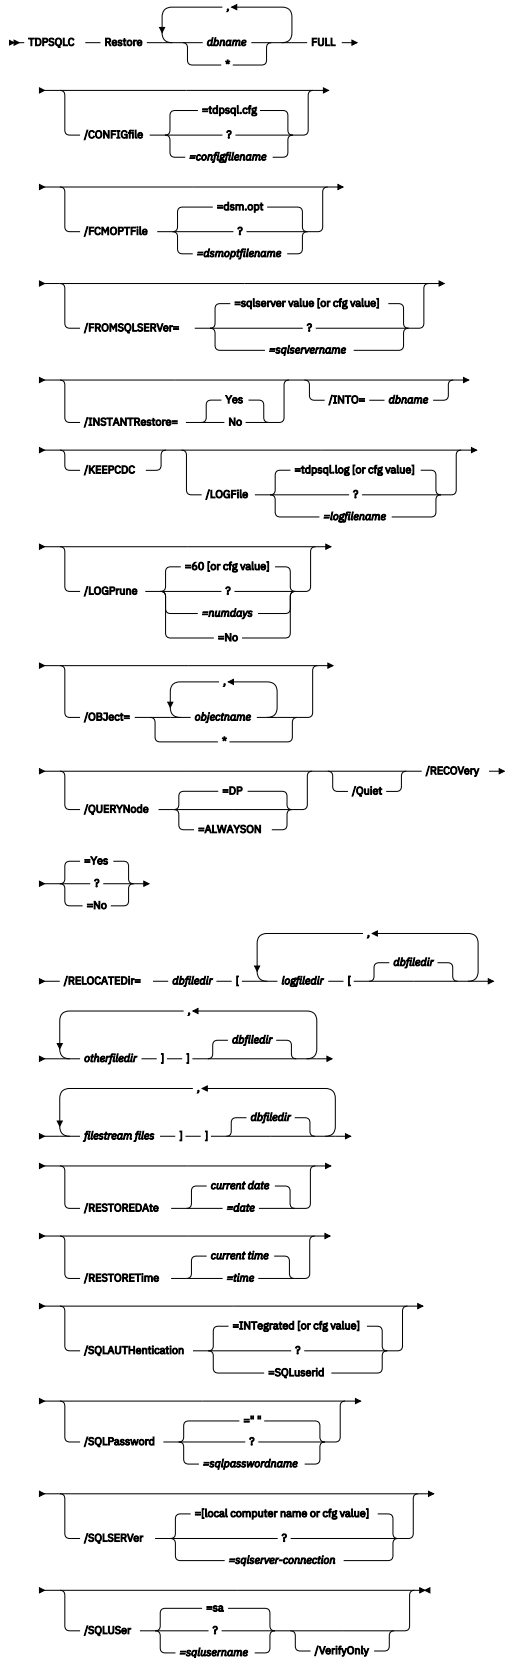
2. SQL Enterprise Manager, SQL Server Application Client, and other SQL Server services can be users of databases and the SQL Server.
- The user that is used by IBM Spectrum Protect Snapshot to log on to the SQL Server must have the SQL Server SYSADMIN fixed server role.
 - You can use the TRANSACT-SQL database consistency checker statement DBCC CHECKDB ('DBNAME') to verify the integrity of the restored SQL databases.

Restore syntax

Use the **restore** command syntax diagrams as a reference to view available options and truncation requirements.

Syntax

TDPSQLC command



Restore positional parameters

Positional parameters immediately follow the **restore** command and precede the optional parameters.

FULL

This option restores all full database backup objects for the SQL databases that you specify.

Restore optional parameters

Optional parameters are used with the **restore** command and positional parameters.

The following are detailed descriptions of each of the optional parameters:

/CONFIGfile=configfilename

The **/CONFIGfile** parameter specifies the name of the IBM Spectrum Protect Snapshot configuration file, which contains the values for the IBM Spectrum Protect Snapshot configurable options. See [“Set command”](#) on page 306 for details on the contents of the file. Considerations:

- configfilename can include a fully qualified path. If configfilename does not include a path, it uses the directory where IBM Spectrum Protect Snapshot is installed.
- If configfilename includes spaces, place it in double quotation marks.
- If you do not specify **/CONFIGfile**, the default value is tdpsql.cfg.
- If you specify **/CONFIGfile** but not configfilename, the default value tdpsql.cfg is used.

/FCMOPTfile=dsmoptfilename

The **/FCMOPTfile** parameter specifies the IBM Spectrum Protect Snapshot options file to use. Considerations:

- The *dsmoptfilename* variable can include a fully qualified path. If you do not include a path, the IBM Spectrum Protect Snapshot installation directory is used.
- If the *dsmoptfilename* variable spaces, enclose it in double quotation marks.
- If you do not specify **/FCMOPTfile**, the default value is dsm.opt.
- If you specify **/FCMOPTfile** but not dsmoptfilename, the default is also dsm.opt.

/FROMSQLSERVer=sqlservername

For **restore**, the **/fromsqlserver** parameter specifies the SQL server that backup objects were backed up from. This parameter is necessary only when the name of the SQL Server to restore to, as determined by the **/sqlserver** parameter, is different from the name of the SQL Server that the backup objects were created from. Use **/fromsqlserver** for **query FCM** commands, but use **/sqlserver** for **query SQL** commands. The default value is the **/sqlserver** value or the value that is set in the IBM Spectrum Protect Snapshot configuration file. If the two SQL Server names are different, you must use this parameter even if **/fromsqlserver** was a non-clustered default instance.

/INSTANTRestore=Yes|No

Use the **/INSTANTRestore** parameter to specify whether to use volume level snapshot or file level copy to restore a VSS backup that is stored on local shadow volumes. An IBM Systems Storage SAN Volume Controller, DS8000, the XIV system, and IBM Storwize V7000 storage subsystem is required to run VSS instant restores.

You can specify:

Yes

Use volume level snapshot restore for a VSS backup that is stored on local shadow volumes if the backup exists on volumes that support it. This option is the default.

No

Use file-level copy to restore the files from a VSS backup that is stored on local shadow volumes. Bypassing volume-level copy means that SQL database files and log files are the only data overwritten on the source volumes.

When you are running VSS instant restore on DS8000 and Storwize family, ensure that any previous background copies that involve the volumes you are restoring, complete before you initiate the VSS instant restore operation.

/INTO=dbname

For **restore** operations, **/INTO** specifies the SQL Server database that you want a backup object that is restored into. This parameter is necessary only when the name of the SQL Server database to restore into is different from the backup object database name. Considerations:

- When you specify **/INTO**, the asterisk (*) wildcard character might not be used in either the command *dbname* variable or the **/INTO dbname** variable.
- There must be exactly one item in the **/INTO dbname** variable list in addition to in the command *dbname* list.
- Make sure to use the **/relocatedir** parameter when you specify **/INTO dbname**.

/KEEPCdc

If change data capture is enabled for an SQL Server database and database tables, change-data-capture records that recorded activity, such as insertions, deletions, and edits to the database tables can be retained when you restore the database. The **/KEEPCdc** parameter is necessary if you are restoring the database to a different database name on the same SQL Server instance or to a different SQL Server instance. When you run restore operations with the **/KEEPCdc** parameter, you must also set **/RECOVERY=Yes**.

This parameter applies to the following types of legacy backups: full, copy-only full, differential, file, set, and group. For VSS restore operations, you do not need to maintain change data capture information.

Note: If you are restoring the database to its original location, that is the SQL Server instance and database name, change data capture records are retained automatically.

For more information about change data capture, see Microsoft documentation.

/LOGFile=logfilename

The **/LOGFile** parameter specifies the name of the activity log that is generated by IBM Spectrum Protect Snapshot. This activity log records significant events such as completed commands and error messages. The IBM Spectrum Protect Snapshot activity log is distinct from the SQL Server error log. The **/LOGFile=** variable identifies the name to be used for the activity log generated by IBM Spectrum Protect Snapshot. Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully qualified path; however, if you specify no path, the file is written to the directory where IBM Spectrum Protect Snapshot is installed.
- You cannot turn off IBM Spectrum Protect Snapshot logging activity. If you do not specify **/LOGFile**, log records are written to the default log file. The default log file is *tdpsql.log*.
- When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot for operations, use the **/LOGFile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OBJECT=*|objectname,...

For restore and deactivate operations, **/OBJECT** specifies that only particular backup objects for the specified SQL databases and backup object type if specified are restored. For query operations, **/OBJECT** includes particular objects and object types in the display. The *objectname* variable specifies the names of the backup objects you want to restore or deactivate. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot. Use **query** to view the names of backup objects. Considerations:

- If you do not specify restore, only the active backup object is included in the restore.
- You can specify the asterisk (*) wildcard character in objectname to replace zero or more characters for each occurrence. Specifying only the wildcard character indicates all backup objects of the specified SQL databases and backup object type.

/QUERYNode=DP|ALWAYSON

Specify whether you want to query standard databases from SQL Server 2012 that were backed up from a standard IBM Spectrum Protect Snapshot for Microsoft SQL Server node or the AlwaysOn node. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node. The default value is DP. To query backups of AlwaysOn Availability databases, specify **/QUERYNode = ALWAYSON**.

/Quiet

The **/Quiet** parameter omits displaying status information from the command. However, the information is appended to the IBM Spectrum Protect Snapshot activity log.

/RECOVery=Yes|No

For restore operations, **/RECOVery** specifies whether you want to restore more to an SQL database that is not on a standby SQL Server. A restored database cannot be used until the **/RECOVery=yes** parameter is administered to the database. You can specify:

Yes (default)

Use when you make a sequence of restores to an SQL database and the current restore is the final in the sequence. Also, use this option when the restore operation is the only restore operation to an SQL database. This option informs the SQL Server that the restore is complete and ready for incompleting transactions to be rolled back.

No

Whenever you make a sequence of restores to an SQL database and the current restore is not the final restore in the sequence.

Not specifying this option automatically rolls back incomplete transactions for the database.

IBM Spectrum Protect Snapshot sorts the restore objects by database name, and, within database name, by backup time from earliest to latest time. A **query FCM** command also displays this order.

/RELOCATEDir=dbfiledir [,logfiledir [,otherfiledir] [,filestream files]]

The **/RELOCATEDir** parameter specifies the new destination locations in which to restore the backed up SQL databases, logs, and SQL Server full-text index files. FILESTREAM files are included for SQL Server 2008 or later versions.

The *dbfiledir* variable specifies the directory location of the SQL database you want to relocate. Note, if the *logfiledir* and *otherfiledir* variables are not specified, the logs and SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *logfiledir* variable specifies the directory location of the SQL log files you want to relocate. Note, if the *logfiledir* variable is not specified, the SQL log files are restored to the directory specified by *dbfiledir*.

The *otherfiledir* variable specifies the directory location of the SQL Server full-text index files you want to relocate. Note, that if the *otherfiledir* variable is not specified, the SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *filestream files* variable specifies the directory location of the SQL Server FILESTREAM data files (SQL Server 2008 or later versions) you want to relocate. Note, if the *filestream files* variable is not specified, the SQL Server FILESTREAM data files are restored to the directory specified by *dbfiledir*. *Filestream files* is available for SQL Server 2008 only.

/RESTOREDate=date

The **/RESTOREDate** parameter specifies a date to which the database identified by *dbname* is to be recovered. The date value must be specified in the same date format that is defined in the IBM Spectrum Protect Snapshot preferences file. If **/RESTOREDate** is not specified but **/RESTORETime** is specified, the **/RESTOREDate** value is the current date. It can be specified only when you restore a full database backup. The **/RESTORETime** parameter cannot be used to restore file, group, and set backups.

/RESTORETime=time

The **/RESTORETime** parameter specifies the time of day to which the database identified by *dbname* is to be recovered. The time value must be specified in the same time format that is defined in the IBM Spectrum Protect Snapshot preferences file. If **/RESTORETime** is not specified but **/RESTOREDate** is specified, the **/RESTORETime** is the current time. It can be specified only when you restore a full database backup. The **/RESTORETime** parameter cannot be used to restore file, group, and set backups.

/SQLAUTHentication=INTegrated | SQLuserid

This parameter specifies the authorization mode that is used when you log on to the SQL Server. The integrated value specifies Windows authentication. The user ID you use to log on to Windows is the same id you use to log on to the SQL Server. This option is the default value. Use the sqluserid value to specify SQL Server user ID authorization. The user ID specified by the /sqluserid parameter is the id that you use to log on to the SQL Server. Any SQL user ID must have the SQL Server SYSADMIN fixed server role.

/SQLPassword=sqlpasswordname

This parameter specifies the SQL password that IBM Spectrum Protect Snapshot uses to log on to the SQL Server that objects are backed up from or restored to. Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user ID for this password must both be configured for SQL Server authentication.
- If you do not specify **/SQLPassword**, the default value is blank (" ").
- If you specify **/SQLPassword** but not **sqlpasswordname**, the default is also blank (" ").

This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/SQLSERVER=sqlserver-connection

The **/SQLSERVER=** parameter specifies the SQL Server that IBM Spectrum Protect Snapshot logs on to. For restore operations, this SQL Server is the one to which backup objects are restored. However, if the backup objects were created from a different SQL Server name, you must use the **/fromsqlserver** parameter. Use **/sqlserver** for the **query SQL** and **backup** commands, but use **/fromsqlserver** for **query FCM** commands.

The *sqlprotocol* variable specifies the communication protocol to use. With this variable, you can specify an *sqlservername*. You can check the SQL connection by using the SQL Server Configuration Manager tool (under SQL Server Native Client Configuration client protocols). You can choose from the following protocols:

<i>Table 41. SQL Server connection protocols</i>		
Protocol Name	Description	Example Usage (with <i>sqlserver-connection</i> details)
lpc	Shared Memory	/sqlserver=lpc:<servername>\<instancename>
np	Named Pipes	/sqlserver=np:<servername>\pipe You can optionally specify a specific named pipe instance. For example, /sqlserver=np: \\hostname\pipe\pipe name By default, the pipe name is <i>sql\query</i> . If you connect to a named instance, the pipe name is typically in the following format: \\<servername>\pipe\MSSQL\$<instancename>\SQL\query
tcp	Transmission Control	/sqlserver=[tcp:]<servername>[\<instancename>][,port]
via	Virtual Interface Adapter	/sqlserver=via:<virtualservername>[\<instancename>]

Note:

- For tcp protocols only, you have the option of defining a *port*. If you do not define a port, the default port value is the SQL default port 1433.
- For the via protocol, SQL Server supports this protocol only through SQL Server 2008 R2.
- To enable Data Protection for SQL Server to communicate with AlwaysOn Availability Group (AAG) instances, it is not possible to connect to the SQL Server using AAG listeners. For backup and restore operations, you must use the local SQL Server instance name (or instance name and port number) to communicate with the AAG. For AAG (or non-AAG instances), you can also specify non-default port numbers.

If you do not specify a protocol, IBM Spectrum Protect Snapshot logs on to the SQL Server according to the first protocol that becomes available.

Considerations:

- The default value is the value that is specified by the SQL Server configurable option in the IBM Spectrum Protect Snapshot configuration file. This value is initially the local computer name.
- If you specify **/sqlserver** but not **sqlservername**, the local computer name is used.
- The following two shortcuts are accepted as the local computer name: . (local) These shortcuts are a period or the word local within parentheses.
- You must specify the name if the SQL Server is not the default instance or is a member of a failover cluster.
- The format of **sqlservername** depends on what type of instance it is and whether it is clustered or not:

Format	Instance?	Clustered?	Name required?
<i>local-computername</i>	default	no	no
<i>local-computername</i> <i>\instancename</i>	named	no	yes
<i>virtualservername</i>	default	yes	yes

Format	Instance?	Clustered?	Name required?
<i>virtualservername</i> <i>\instancename</i>	named	yes	yes

localcomputername

The network computer name of the computer on which the SQL Server and IBM Spectrum Protect Snapshot are stored. The TCP/IP host name might not always be the same.

instancename

The name given to the named instance of the SQL Server that is specified during installation of the instance.

virtualservername

The name that is given to the clustered SQL Server specified during clustering service setup. This name is not the cluster or node name.

/SQLUser=sqlusername

The **/SQLUser** parameter specifies the name that IBM Spectrum Protect Snapshot uses to log on to the SQL Server. Considerations:

- Using this parameter means that you are using SQL Server authentication. The SQL Server and the SQL user ID for this password must both be configured for SQL Server authentication.
- The SQL user ID must have the SQL Server SYSADMIN fixed server role.
- If you do not specify **/SQLUser**, the default is sa.
- If you specify **/SQLUser** but not **sqlusername**, the default is also sa.
- This parameter is ignored if you use the **/sqlauth=integrated** parameter with it.

/VerifyOnly

The **/VerifyOnly** parameter specifies that a database restore action reads backup data to verify the integrity of the data only; it does not save the backup to disk or overwrite the current database of that name on the SQL server. Before you restore a backup, you can use this parameter to evaluate whether the backup volume is complete and can be read.

If this parameter is not specified, the restore action both verifies the integrity of the backup and also saves the backup to disk on the SQL server.

Restriction: The **/VerifyOnly** parameter is available only for legacy database backups. This parameter is only a command optional parameter, and it cannot be set as a configuration option.

To verify a backup done with multiple stripes, set the **Verify Only** parameter to **Yes** and ensure that the number of stripes that are used for the restore is equal to or greater than the number of stripes that are used for the backup you are verifying. If it is not, the **Restore VerifyOnly** operation terminates with an error.

Restore output examples

These output examples provide a sample of the text, messages, and process status that displays when you use the **restore** command.

In this example, the **tdpsqlc restore db1 full** command restores a full backup of database *db1*. The following output is displayed:

```
IBM FlashCopy Manager for Databases:
FlashCopy Manager for Microsoft SQL Server
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.
```

```
Connecting to SQL Server, please wait...
```

```
Querying Virtual Server for Backups ....
```

```
Starting Sql database restore...
```

```
Beginning VSS restore of 'db1'...
```

```
Files Examined/Completed/Failed: [ 3 / 3 / 0 ] Total Bytes: 6029825
```

```
VSS Restore operation completed with rc = 0
```

```
Files Examined : 3
```

```
Files Completed : 3
```

```
Files Failed : 0
```

```
Total Bytes : 6029825
```

Restorefiles command

Use the **restorefiles** command to restore VSS-based backups on the IBM Spectrum Protect (/BACKUPDESTINATION=TSM), or stored locally (/BACKUPDESTINATION=LOCAL).


Considerations

- The **restorefiles** command restores .mdf, ldf, and other flat files from a specified IBM Spectrum Protect VSS-based backup into a specified directory.
- A destination directory can be specified as a directory on a fixed file system such as C:\temp, or on a network share that is accessible to the IBM Spectrum Protect Snapshot Remote Agent (VSS Requestor). It is not possible to use a mapped network drive as a destination directory.
- The **restorefiles** command does not restore the data to the SQL Server.
- This command does not require the SQL Server to be installed on the system where the **restorefiles** command is run.
- A restore continues until it is completed unless the destination volume does not have enough space to fulfill the restore operation.
- VSS-based backups that are on the IBM Spectrum Protect Snapshot (/BACKUPDESTINATION=TSM) can be restored by using **restorefiles** on the same system that ran the VSS-based backup, or by running the command on a system that installed and configured the IBM Spectrum Protect Snapshot client.
- The directory that is specified in the **restorefiles** command appended the VSS component name so that multiple databases can be restored to the same target directory.
- VSS-based backups that are stored on the local system by using a persistent snapshot (/BACKUPDESTINATION=LOCAL), can be restored only by running the **restorefiles** command on the same system that ran the VSS-based backup, and has access to the persistent snapshot.
- To run a full restore: `tdpsqlc restorefiles DBName1 FULL /backupmethod=vss /relocatedir=d:\temprestore`
- Use /RELOCATEDIR to specify the destination directory for the flat files. If this option is not specified, the destination directory defaults to the current working directory.
- If you are in a non-clustered environment, you can restore only a local snapshot to the system that generated the snapshot. Or for cluster environments, you can run a **restorefiles** command from any of the systems in the cluster.

Restorefiles syntax

Use the **restorefiles** command syntax diagram as a reference for available options and truncation requirements.

TDPSQLC command



Restorefiles positional parameters

Positional parameters immediately follow the **restorefiles** command and precede the optional parameters.

The following positional parameters specify the object to restore:

TDPSQLC * | *componentname1*, ..., *componentnameN* FULL

*

Sequentially restore all flat files for the database.

The following positional parameter specifies the type of backup from which the files are restored:

FULL

Restore the files from a full type backup for VSS.

Restorefiles optional parameters

Optional parameters follow the **restorefiles** command and positional parameters.

/BACKUPDESTINATION

VSS backups that are on the IBM Spectrum Protect server are restored by using the **restorefiles** command with **/BACKUPDESTINATION=TSM**. VSS backups that are running on a local system that uses a persistent snapshot are restored by using the **restorefiles** command with **/BACKUPDESTINATION=LOCAL**. TSM is the default destination for **restorefiles**.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name of the IBM Spectrum Protect Snapshot configuration file that contains the values for the IBM Spectrum Protect Snapshot configuration options.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpsql.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\file.cfg"
```

/FROMSQLserver=*sqlservername*

Use the **/FROMSQLserver** parameter to specify the name of the SQL Server where the original backup was completed. The default is the local SQL Server name. To restore availability databases, specify the AlwaysOn Availability group.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\mytdpsqlserver.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, *tdpsqlserver.log*.

The **/logfile** parameter cannot be turned off, logging always occurs.

When you use multiple simultaneous instances of IBM Spectrum Protect Snapshot to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This function directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option *No* can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTWait=Yes|No

The **/MOUNTWait** parameter is used to specify whether IBM Spectrum Protect Snapshot waits for removable media to mount, such as tapes or CDs, or stops the operation. This situation occurs when the IBM Spectrum Protect Snapshot is configured to store backup data on removable media and waits for a required storage volume to be mounted. This parameter is not valid for all backup types; the parameter does not work with DIFFFULL or LOG backup types.

You can specify these options:

Yes

Wait for tape mounts. This option is the default.

No

Do not wait for tape mounts.

/OBJECT=object name

Use the **/object** parameter to specify the name of the backup object files that you want to restore. The object name uniquely identifies each backup object and is created by IBM Spectrum Protect Snapshot.

Use the IBM Spectrum Protect Snapshot **query tsm** command to view the names of the backup objects.

/Quiet

This parameter prevents status information from being displayed. This function does not affect the level of information that is written to the activity log.

/RELOCATEDir=dbfiledir[,logfiledir [,otherfiledir] [,filestream files]]

The **/relocatedir** parameter specifies the destination locations in which to restore the flat files. This restore includes databases, logs, and FILESTREAM files. It is not possible to use a mapped network drive as a destination directory.

The *dbfiledir* variable specifies the directory location of the SQL database you want to relocate. If the *logfiledir* or *otherfiledir* variables are not specified, the logs and SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The *logfiledir* variable specifies the directory location of the SQL log files you want to relocate. If the *logfiledir* variable is not specified, the SQL log files are restored to the directory specified by *dbfiledir*.

The *otherfiledir* variable specifies the directory location of the SQL Server full-text index files you want to relocate. If the *otherfiledir* variable is not specified, the SQL Server full-text index files are restored to the directory specified by *dbfiledir*.

The **restorefiles** operation creates a subdirectory under the root directory that contains the name of the database name. Restored files are placed in that subdirectory. If the **/relocatedir** parameter is not specified, the files are restored into the directory where the **restorefiles** command is issued. For example, if IBM Spectrum Protect Snapshot is installed in the `c:\Program Files\Tivoli\TSM\TDPSQLC` directory and the following command is issued from `E:\Somedir`:

```
e:\Somedir> c:\"Program Files"\Tivoli\TSM\TDPSQLC\tdpsqlc restorefiles
db1 full
```

then the files are restored to the subdirectories in the `e:\Somedir` location:

```
e:\Somedir\db1\db1.mdf
e:\Somedir\db1\db1.ldf
```

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on. You can store the node name in the IBM Spectrum Protect options file (`dsm.opt`). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect Snapshot options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/TSMPassword=tsmpassword

. Use the *tsmpassword* variable to refer to the password that IBM Spectrum Protect Snapshot uses to log on to IBM Spectrum Protect. If you specified `PASSWORDACCESS GENERATE` in the options file (`dsm.opt`), supplying the password is not necessary because the one that is stored in the registry is used. Store the password in the registry by specifying the IBM Spectrum Protect password for the first connection.

If you do specify a password with this parameter when `PASSWORDACCESS GENERATE` is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If `PASSWORDACCESS PROMPT` is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The password can be up to 63 characters in length.

Restorefiles examples

This output example provides a sample of the text, messages, and process status that displays when you use the **restorefiles** command.

This command, `tdpsqlc restorefiles Finance FULL /backupdestination=local /RELOCATEDir=e:\test /FROMSQLServer=sqlsrv12`, restores VSS files from a FULL type backup of the *Finance* database from the SQL Server named *sqlsrv12* into the `e:\test` directory. The restored files are displayed:

```
e:\test\Finance\finance.mdf  
e:\test\Finance\finance_log.ldf
```

Set command

Use the **set** command to set the IBM Spectrum Protect Snapshot for SQL Server configuration parameters that are defined in the IBM Spectrum Protect Snapshot for SQL Server configuration file, `tdpsql.cfg` by default.

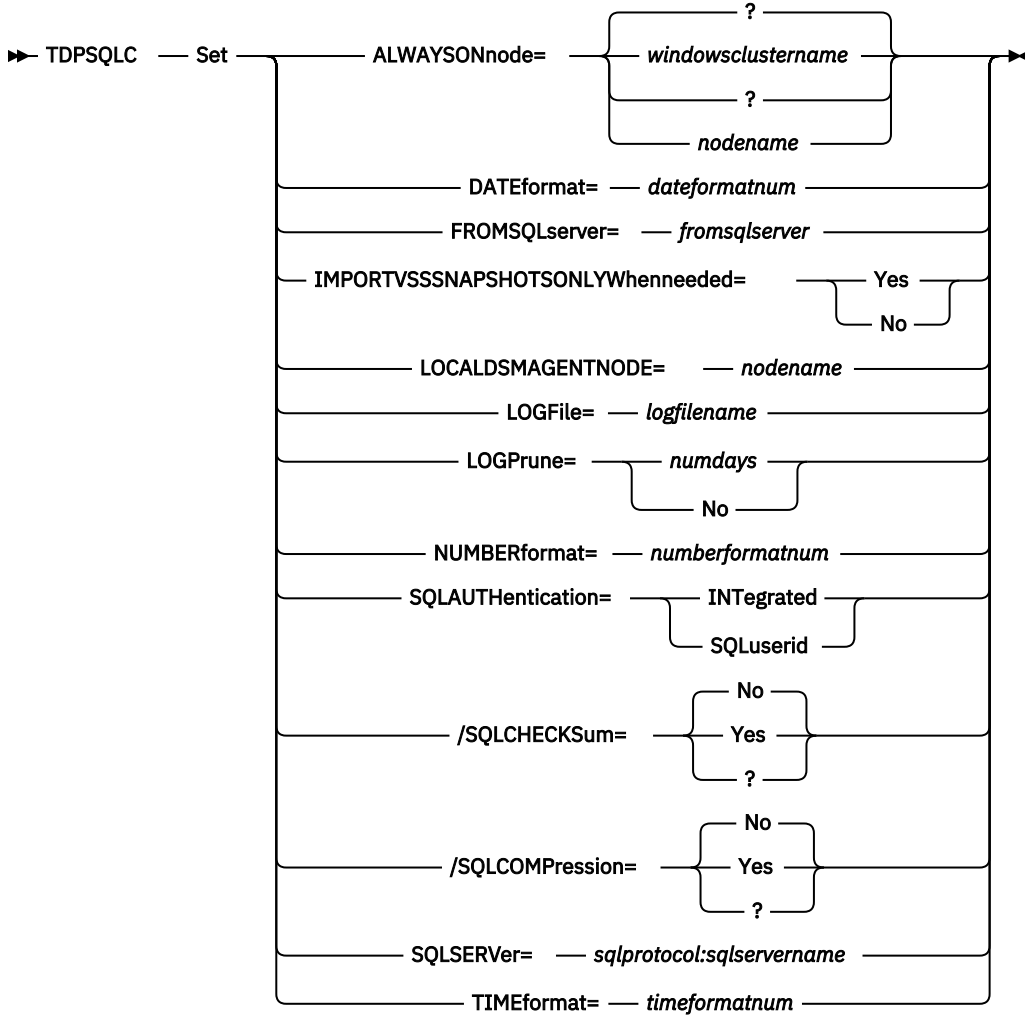
Use the **set** command to change the values for the IBM Spectrum Protect Snapshot configurable parameters and options. The values are saved in a configuration file. The default file is `tdpsql.cfg`. Configuration values can also be set from the **Edit** menu in the GUI.

Note: If a configuration file is not specified, the `tdpsql.cfg` values are used, and a default configuration file is created with just the *lastprunedate* value. If an invalid or non-existent file is specified, the default values are used.

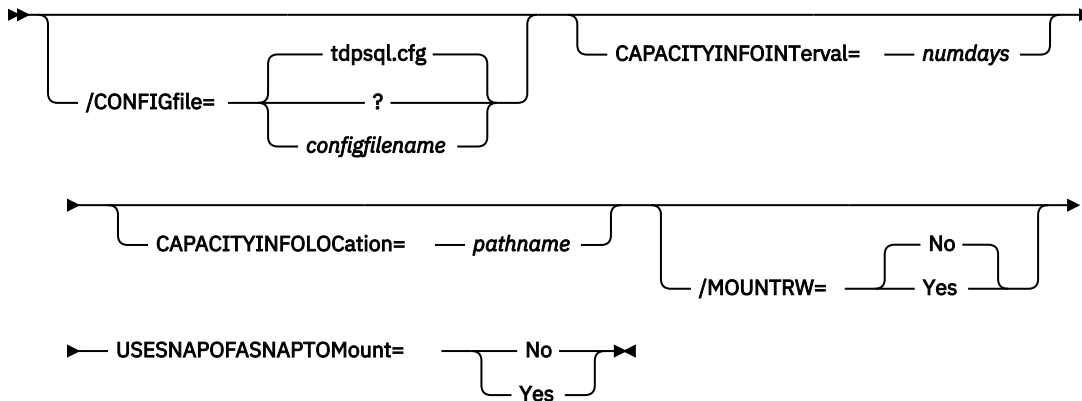
Set syntax

Use the **set** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command



Set Optional Parameters



Set positional parameters

Positional parameters immediately follow the **set** command and precede the optional parameters.

To set default values in the IBM Spectrum Protect Snapshot configuration file, specify one of the following when you issue a **set** command.

ALWAYSONNode=nodename

Specify the IBM Spectrum Protect node name that is used to back up AlwaysOn availability databases with SQL Server 2012 and later versions. This parameter is required when you are configuring IBM Spectrum Protect Snapshot with SQL Server 2012 and later versions. All availability databases in an availability group are backed up under this node name, regardless of which availability replica they are from. The databases that are not in an availability group are backed up under the standard IBM Spectrum Protect Snapshot node name unless you specify the **USEALWAYSONnode** parameter.

CAPACITYINFOInterval=numdays

Use the **CAPACITYINFOInterval** positional parameter to specify how often you want the capacity metrics report to be generated. The report, in an XML file format, is generated automatically at the end of a backup operation. The valid value range is 1 - 365 and the default value is 7 days, which means the report is generated once every 7 days.

CAPACITYINFOLOcation=pathname

Use the **CAPACITYINFOLOcation** positional parameter to specify the location where the capacity metrics report is to be created. If you do not specify a location, the report is not generated.

DATEformat=dateformatnum

The **DATEformat** parameter selects the format that you want to use to display dates.

The *dateformatnum* variable can range 1 - 7. The initial value is 1. The number values specify the following formats:

- 1** MM/DD/YYYY.
- 2** DD-MM-YYYY.
- 3** YYYY-MM-DD.
- 4** DD.MM.YYYY.
- 5** YYYY.MM.DD.
- 6** YYYY/MM/DD.
- 7** DD/MM/YYYY.

Changes to the value of the **dateformat** parameter can result in an undesired pruning of the IBM Spectrum Protect Snapshot log file (*tdpsql1.log* by default). You can avoid losing existing log file data by running one of the following tasks:

- After you change the value of the **dateformat** parameter, make a copy of the existing log file before you run IBM Spectrum Protect Snapshot.
- Specify a new log file with the **/logfile** parameter.

FROMSQLSERVER=sqlservername

The **fromsqlserver** parameter specifies the SQL Server that backup objects were backed up from. This parameter is necessary only when the name of the SQL Server to restore to, as determined by the **sqlserver** parameter, is different from the name of the SQL Server that the backup objects were created from. Use **fromsqlserver** for **query FCM**, but use **sqlserver** for **query SQL** commands. The default value is the *sqlserver* value or the value that is set in the IBM Spectrum Protect Snapshot configuration file.

IMPORTVSSSNAPSHOTSONLYWhenneeded

Use the **/IMPORTVSSSNAPSHOTSONLYWhenneeded** parameter to specify whether IBM Spectrum Protect Snapshot automatically imports VSS snapshots to the Windows system where the snapshots are created.

Specify one of the following values:

Yes

Import VSS snapshots to the Windows system where the snapshots are created. The option is the default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required.

Tip: For instant restore processing on IBM and non-IBM storage systems, you must specify the Yes option to enable the storage system to create transportable snapshots during backups.

No

Do not create transportable VSS snapshots during backup processing, and do not automatically import the snapshot to storage systems after the backup is completed.

LOCALDSMAgentnode=nodename

Specify the node name of the local system that runs the VSS backups. This positional parameter must be specified for VSS operations to be done.

LOGFile=logfilename

The **logfile** parameter specifies the name of the activity log that is generated by IBM Spectrum Protect Snapshot. The activity log records significant events such as completed commands and error messages. This log is distinct from the SQL Server error log. The *logfilename* variable identifies the name to be used for the activity log generated by IBM Spectrum Protect Snapshot.

Considerations:

- If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.
- The file name can include a fully qualified path; however, if you specify no path, the file is written to the directory where IBM Spectrum Protect Snapshot is installed.
- You cannot turn off IBM Spectrum Protect Snapshot logging activity. If you do not specify **/logfile**, log records are written to the default log file. The default log file is `tdpsql.log`.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, some days of data are saved. By default, 60 days of log entries are saved. The option No can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify no, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

NUMBERformat=numberformatnum

The **numberformat** parameter specifies the format of the numbers that are displayed by IBM Spectrum Protect Snapshot. The *numberformatnum* variable can range 1 - 6. The initial value is 1. The number values specify the following formats:

- 1** 1,000.00
- 2** 1,000,00
- 3** 1 000,00
- 4** 1 000.00
- 5** 1.000,00
- 6** 1'000,00

SQLAUTHentication=INTEgrated | SQLuserid

This parameter specifies the authorization mode that is used when you log on to the SQL Server. The integrated value specifies Windows authentication. The user ID that you use to log on to Windows is the same ID you use to log on to the SQL Server. This option is the default value. Use the *sqluserid* value to specify SQL Server user ID authorization. The user ID specified by the *sqluserid* parameter is the ID you use to log on to the SQL Server. That user ID must have the SQL Server SYSADMIN fixed server role.

SQLSERVer=sqlprotocol:sqlservername

The **SQLSERVersqlserver** parameter specifies the SQL Server that IBM Spectrum Protect Snapshot logs on to. This SQL Server is the one that backup objects are restored to. However, if the backup objects were created from a different SQL Server name, you must use the **fromsqlserver** parameter. Use **sqlserver** for the **query SQL** command. The *sqlprotocol* variable specifies the communication protocol to use. You can specify one of the following protocols:

- *lpc*: Use Shared Memory protocol.
- *np*: Use Named Pipes protocol.
- *tcp*: Use Transmission Control protocol.
- *via*: Use Virtual Interface Architecture protocol.

If no protocol is specified, IBM Spectrum Protect Snapshot logs on to the SQL Server according to the first protocol that becomes available.

TIMEformat=timeformatnum

The **timeformat** parameter specifies the format of the times that are displayed by IBM Spectrum Protect Snapshot. The *timeformatnum* variable can range 1 - 4. The initial value is 1. The number values specify the following formats:

- 1** 23:00:00
- 2** 23,00,00
- 3** 23.00.00
- 4** 11:00:00A/P

Changes to the value of the **timeformat** parameter can result in an undesired pruning of the IBM Spectrum Protect Snapshot log file (*tdpsql.log* by default). You can avoid losing existing log file data by doing one of the following choices:

- After you change the value of the **timeformat** parameter, make a copy of the existing log file before you run IBM Spectrum Protect Snapshot.
- Specify a new log file with the **/logfile** parameter.

USEALWAYSONnode

Specify this parameter to back up standard databases on SQL Server 2012 and later versions by using the AlwaysOn node. By setting this parameter, you can back up all availability databases and standard databases under a single node to help you to manage your database backups more easily. By default, SQL Server 2012 and later version availability databases are backed up to the AlwaysOn node.

Set optional parameters

Optional parameters follow the **set** command and positional parameters.

/CONFIGfile=configfilename

The **/configfile** parameter specifies the name of the IBM Spectrum Protect Snapshot configuration file, which contains the values for the IBM Spectrum Protect Snapshot configurable options.

Considerations:

- *configfilename* can include a fully qualified path. If *configfilename* does not include a path, it uses the directory where IBM Spectrum Protect Snapshot is installed.
- If *configfilename* includes spaces, place it in double quotation marks.
- If you do not specify **configfile**, the default value is `tdpsql.cfg`.
- If you specify **configfile** but not *configfilename*, the default value `tdpsql.cfg` is used.

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

Perform a read-only mount operation.

Yes

Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPTOMount** parameter in the configuration file.

- If **USESNAPOFASNAPTOMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
- If **USESNAPOFASNAPTOMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which require IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which require IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate more target volumes on your storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

SQLCHECKSum=Yes | No

The **SQLCHECKSum** parameter specifies whether SQL checksum processing is set for all legacy database backups. If you do not specify a value of **Yes** for this parameter, the value that is specified in the `tdpsql.cfg` file is used. If no value is specified in the `tdpsql.cfg` file, the default value of **No** is used.

The **SQLCHECKSum** parameter is only available with legacy backups.

/SQLCOMPression=Yes | No

The **/sqlcompression** parameter specifies whether SQL compression is applied. If you do not specify **/sqlcompression**, the default value **No** is used.

This parameter is only applicable on systems that run SQL Server 2008 or later versions. For SQL Server 2008, you can run backup compression only on Enterprise Edition. For SQL Server 2008 R2, backup compression is supported on Standard, Enterprise, and Datacenter editions.

Set output example

These output examples provide a sample of the text, messages, and process status that displays when you use the **set** command.

The following specifies the *mutalisk* server as the default SQL server in the configuration file.

Command:

```
tdpsqlc set sqlserver=mutalisk
```

Output:

```
FMY5054I The configuration option was set successfully.
```


Unmount Backup command

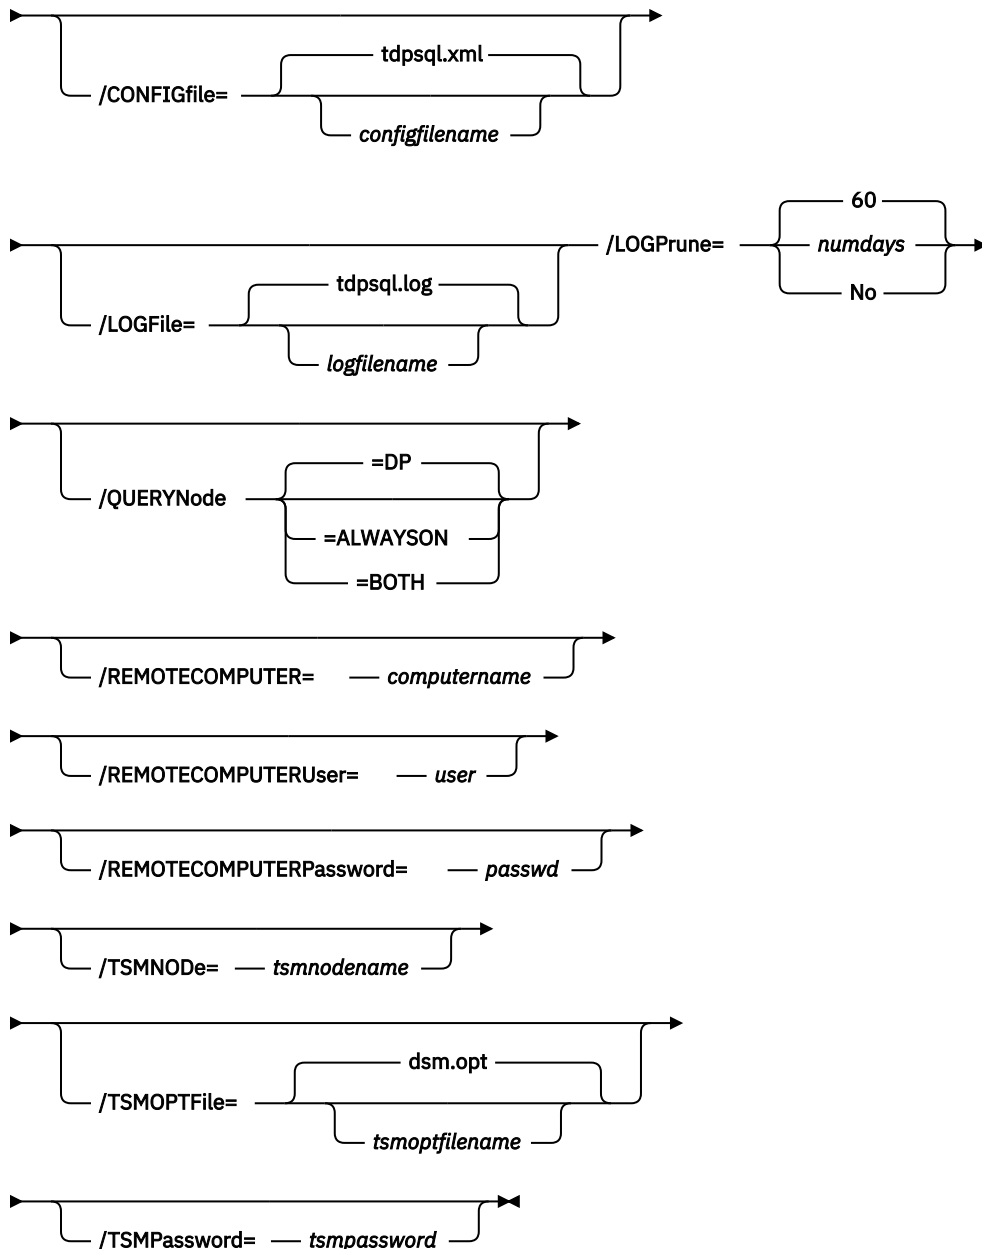
Use the **unmount backup** command to unmount backups that are previously mounted, and are managed by IBM Spectrum Protect Snapshot for SQL Server.

Unmount Backup syntax

Use the **unmount backup** command syntax diagrams as a reference to view available options and truncation requirements.

TDPSQLC command

➤ TDPSQLC — UNMOUNT BACKUP — *mount point root directory* ➤



Unmount Backup positional parameter

The positional parameter immediately follows the **unmount backup** command and precedes the optional parameters.

mount points root directory

Unmount Backup optional parameters

Optional parameters follow the **unmount backup** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the configuration file that contains the values to use for an **unmount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `tdpsql.cfg`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\tdpsql.cfg"
```

/LOGFile=*logfile*

Use the **/logfile** parameter to specify the name of the activity log file. The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\tdpsql.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `tdpsql.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=*numdays***|No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/QUERYNode=DP | ALWAYSON | BOTH

Specify whether you want to query standard databases from SQL Server 2012 and later database versions that were backed up from a standard Data Protection for SQL Server node, the AlwaysOn node, or both nodes. This parameter is ignored for availability databases because the availability databases are always backed up under the AlwaysOn node.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to unmount the data.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\dsm.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Unmount backup example

An example of how to use the **UNMOUNT BACKUp** command is provided.

```
TDPSQLC UNMOUNT BACKUp K:\MP-dir
```

IBM Spectrum Protect Snapshot commands for custom applications and file systems

The name of the IBM Spectrum Protect Snapshot for custom applications and file systems command-line interface is **fcmccli.exe**. By default, this program is in the IBM Spectrum Protect Snapshot installation directory (C:\Program Files\Tivoli\FlashCopyManager\).

Command-line parameter characteristics

Review these parameter characteristics before you attempt a command-line operation.

- Positional parameters do not include a leading slash (/) or dash (-)
- Optional parameters can display in any order after the required parameters
- Optional parameters begin with a forward slash (/) or a dash (-)
- Minimum abbreviations for keywords are indicated in uppercase text
- Some keyword parameters require a value
- For those keyword parameters that require a value, the value is separated from the keyword with an equal sign (=)
- If a parameter requires more than one value after the equal sign, the values are separated with commas
- Each parameter is separated from the others by using spaces
- If a parameter's value includes spaces, the value must be enclosed in double quotation marks
- A positional parameter can display only one time per command invocation

For help in reading syntax diagrams, see [“Reading syntax diagrams”](#) on page xi.

Command-line interface help

Issue the `fcmccli ?` or `fcmccli help` command to display help for the command-line interface.

Related tasks

[Protecting custom application and file system data](#)

With IBM Spectrum Protect Snapshot, you can back up and restore custom application and file system data, and protect your environment.

Backup command

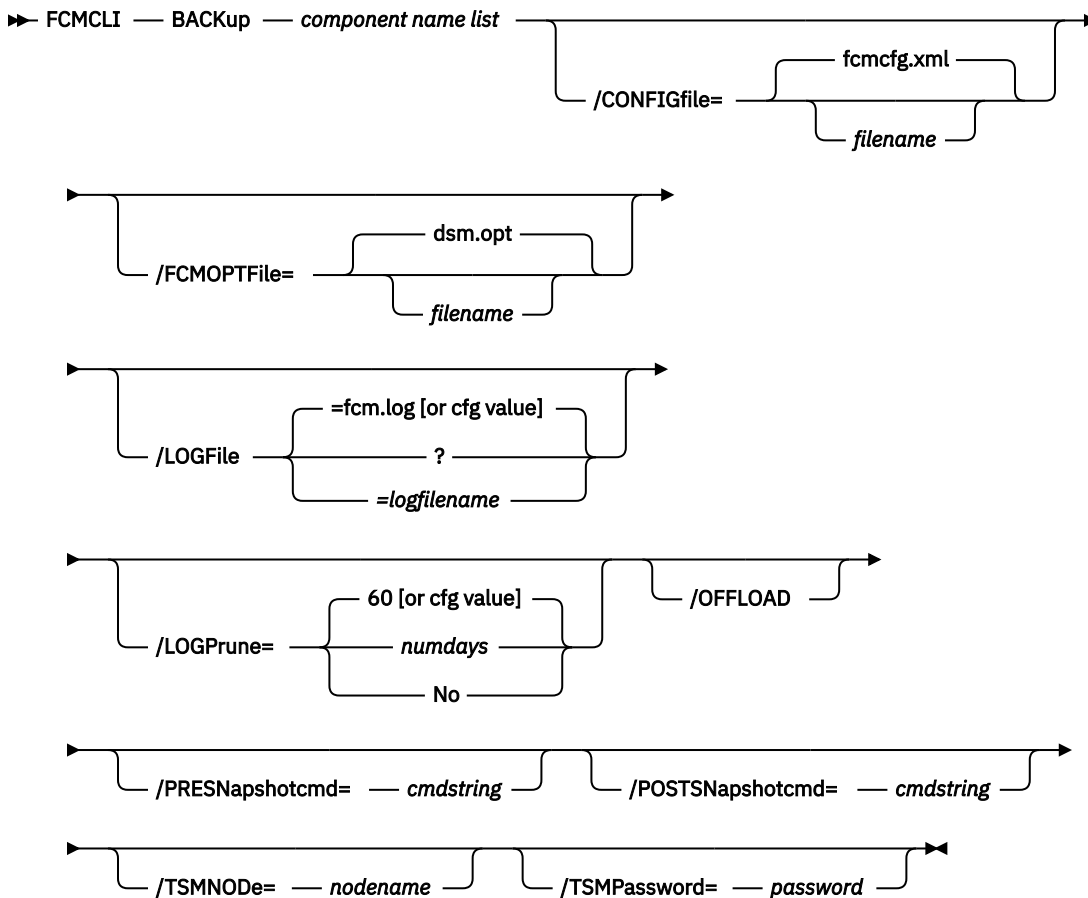
Use the **backup** command to create a VSS snapshot backup of volumes and mount points to local shadow volumes.

The VSS snapshot is managed by IBM Spectrum Protect Snapshot or IBM Spectrum Protect.

Backup syntax

Use the **backup** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command



Backup positional parameter

The positional parameter immediately follows the **backup** command and precedes the optional parameters.

Specify the following positional parameter with the **backup** command:

component name list

Specify a list, of volume or mount points that are separated by commas to back up.

Backup optional parameters

Optional parameters follow the **backup** command and positional parameters.

/CONFIGfile=filename

Use the **/CONFIGfile** parameter to specify the name (*filename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **backup** operation.

The *filename* variable can include a fully qualified path. If the *filename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *filename* variable is not specified, the default value is `fcmcfg.xml`.

If the *filename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/FCMOPTFile=filename

Use the *filename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *filename* variable includes spaces, enclose the entire **/FCMOPTFile** parameter entry in double quotation marks. For example:

```
/FCMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/LOGFile=filename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *filename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *filename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *filename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `fcm.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/OFFLOAD

Specify this option if, after the VSS snapshot is complete, you want to offload the transfer of the data from the IBM Spectrum Protect server to the system specified by the **REMOTEDSMAGENTNODE**

parameter. This option is only valid when the **BACKUPDESTINATION** parameter is set to either TSM or BOTH. The default is to not offload data.

/PRESNapshotcmd=cmdstring

The ***/PRESNapshotcmd*** parameter runs a command or script before a snapshot operation begins. You can use this optional parameter to quiesce an application before a snapshot is created. You can then restart the application after the snapshot is started by using the ***/POSTSNapshotcmd*** optional parameter. The *cmdstring* variable specifies the command to run before the snapshot operation begins. You must specify the fully qualified path name for the command script.

/POSTSNapshotcmd=cmdstring

The ***/POSTSNapshotcmd*** parameter runs a command or script after a snapshot operation ends. You can use this optional parameter to resume the application after the snapshot is created. This parameter is used with the ***/PRESNapshotcmd*** parameter. The *cmdstring* variable must be a fully qualified path.

/TSMNODE=nodename

Use the *nodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMPassword=password

Use the *password* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Backup examples

These output examples provide a sample of the text, messages, and process status that displays when you use the **backup** command.

In this example, the `backup c:,d:` command is run from the **Automate** tab integrated command line.

The following output is displayed:

```
Preparing for a BACKUP operation, please wait...
Connecting to IBM Spectrum Protect Snapshot Server as node 'MALTA_FS'...
Connecting to Local DSM Agent 'MALTA'...
Starting component backup...
Beginning VSS backup of 'C:', 'D:'...
VSS Backup operation completed with rc = 0.
Elapsed Processing Time: 118.52 seconds
Completed
```

In this example, the backup `c:,d: /PRESNapshotcmd="STOPDB.CMD" /POSTSNapshotcmd="STARTDB.CMD"` is run from the Automate integrated command line. The following output is displayed:

```
C:\Program Files\Tivoli\FlashCopyManager>fccli back c:,d:
/presn="C:\Program Files\Tivoli\FlashCopyManager\stopdb.cmd"
/postsn="C:\Program Files\Tivoli\FlashCopyManager\startdb.cmd"

IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a BACKUP operation, please wait...

Connecting to IBM Spectrum Protect Snapshot Server as node 'MALTA_FS'...
Connecting to Local DSM Agent 'MALTA'...
Starting component backup...

Beginning VSS backup of 'C:', 'D:'...

VSS Backup operation completed with rc = 0.

Elapsed Processing Time: 130.16 seconds
```

Delete backup command

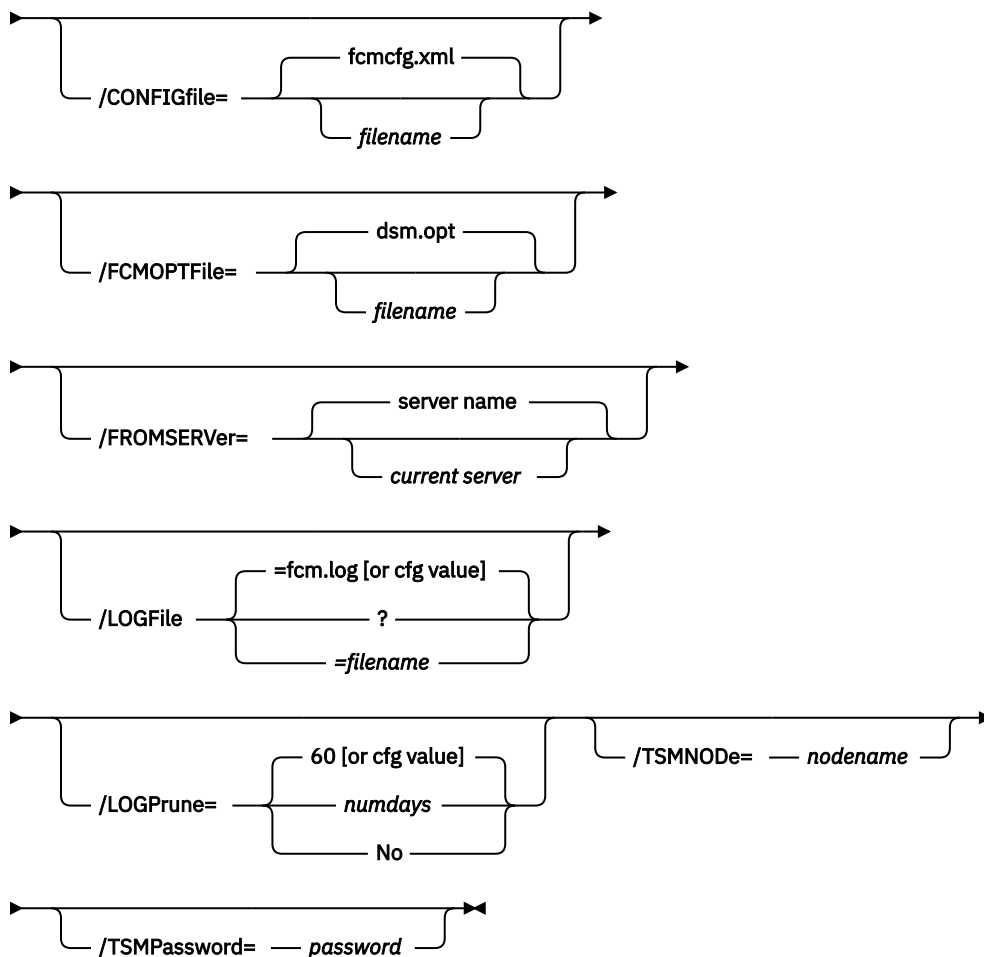
Use the **delete backup** command to delete IBM Spectrum Protect Snapshot snapshot backups from local shadow volumes.

- If you delete multiple LOCAL snapshots that are stored on SAN Volume Controller or Storwize family Space Efficient volumes (SEV), you must do so in the same order in which you created the snapshots. That is, you must delete the oldest one first, followed by the second oldest. Failure to delete them in this order can cause removal of other snapshots of the same source.
- If you mount a local VSS COPY type backup as a snap of a snap, the snap of a snap volume is also deleted along with the VSS backup.

Delete backup syntax

Use the **delete backup** command syntax diagrams as a reference to view available options and truncation requirements.

►► FCMCLI — DELete BACKUp — *component name list* →



Delete backup positional parameter

The positional parameter immediately follows the **delete backup** command and precedes the optional parameters.

Specify the following positional parameter with the **delete backup** command:

component name list

Specify a list of volume or mount points to delete. The list must contain all non-qualified objects or all qualified objects. The list cannot contain a combination of non-qualified objects and qualified objects.

Specify the component name list by using the following syntax:

```
object-1[(object-1-id)][,object-2[(object-2-id)]...]
```

For example:

```
fcmcli delete backup g:(20110311124516),h:(20110211034512),r:(20101114164310)
```

The following example is for a non-qualified object `object-1`:

```
delete backup g:
```

The following example is for a qualified object object-1 (object-1-id):

```
delete backup g:(20110815064316)
```

Use the **query backup** command to find the Object Name identifier.

```
Backups for Volume/Mount Point: 'D:'
=====
Volume/Mount Point      : D:
Volume GUID             : 3487bc7e-4465-11dc-81cc-001a640a19f2
Server                  : MALTA
Volume Occupancy        : 17.40GB
Backup Date/Time        : 08/30/2011 04:07:04
Backup State            : Active
Management Class        : DEFAULT
Mounted as              :
Object Name              : 20110830040704
Instant Restore Supported : No

Completed
```

Delete backup optional parameters

Optional parameters follow the **delete backup** command and positional parameter.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **delete backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `fcmcfg.xml`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/FCMOPTfile=filename

Use the *filename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *filename* variable includes spaces, enclose the entire **/FCMOPTfile** parameter entry in double quotation marks. For example:

```
/FCMOPTfile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/FROMSERVER=server-name

Use the **/fromserver** parameter to specify the name of the server where the original backup was done. The default is the local server.

/LOGFile=logfilename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *fcm.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option *No* can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/TSMNode=nodename

Use the *nodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to **PROMPT**. This parameter is not valid when **PASSWORDACCESS** is set to **GENERATE** in the options file.

/TSMPassword=password

Use the *password* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If `PASSWORDACCESS PROMPT` is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Delete backup example

This output example provides a sample of the text, messages, and process status that displays when you use the `delete backup` command.

In this example, the `fcmlcli delete backup G:,H:` command deletes the backups of volumes G and H. The following output is displayed:

```
Backup(s) to be deleted:
G: and H: : VSS : full : 03/12/2014 12:04:33
VSS Delete backup operation completed with rc = 0
Files Examined      : 2
Files Completed     : 2
Files Failed        : 0
Total Bytes         : 0
```

Help command


Use the `fcmlcli help` command to display help for IBM Spectrum Protect Snapshot commands.

This command lists one or more commands and their parameters. For a language other than English, you might be required to set the width of your screen display. Choose a value greater than 80 characters to view the entire help description in a screen. For example, set the screen width to 100 characters.

Help syntax

Use the `help` command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command



Help positional parameters

Positional parameters follow the IBM Spectrum Protect Snapshot `help` command.

The following positional parameters specify the help to be displayed:

* | *command*

Identifies the specific IBM Spectrum Protect Snapshot command that is to be displayed. If you specify the asterisk (*) wildcard character, help for all IBM Spectrum Protect Snapshot commands is displayed.

* | *subcommand*

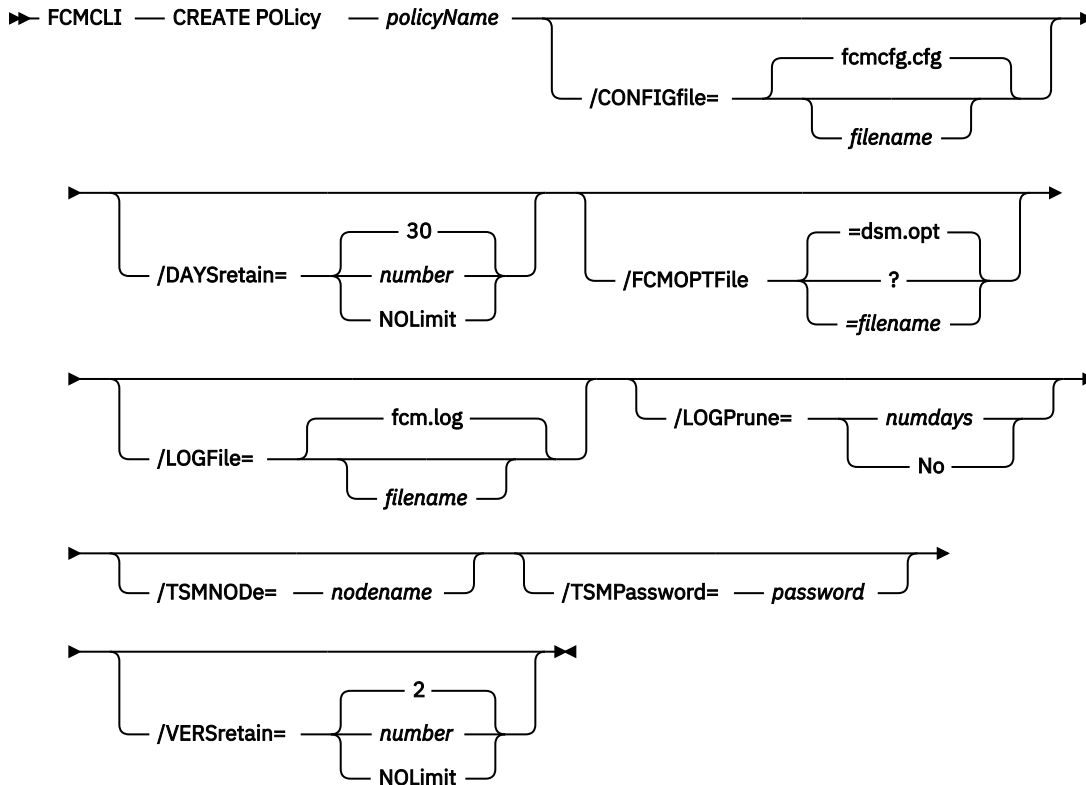
Help can be displayed for commands that have several subcommands, for example, the `query components` command. If you do not specify a subcommand or asterisk (*) wildcard character, help for all IBM Spectrum Protect Snapshot `query components` commands is displayed.

Policy commands for IBM Spectrum Protect Snapshot

Create Policy

This command is used to create a policy.

FCMCLI command



Parameters:

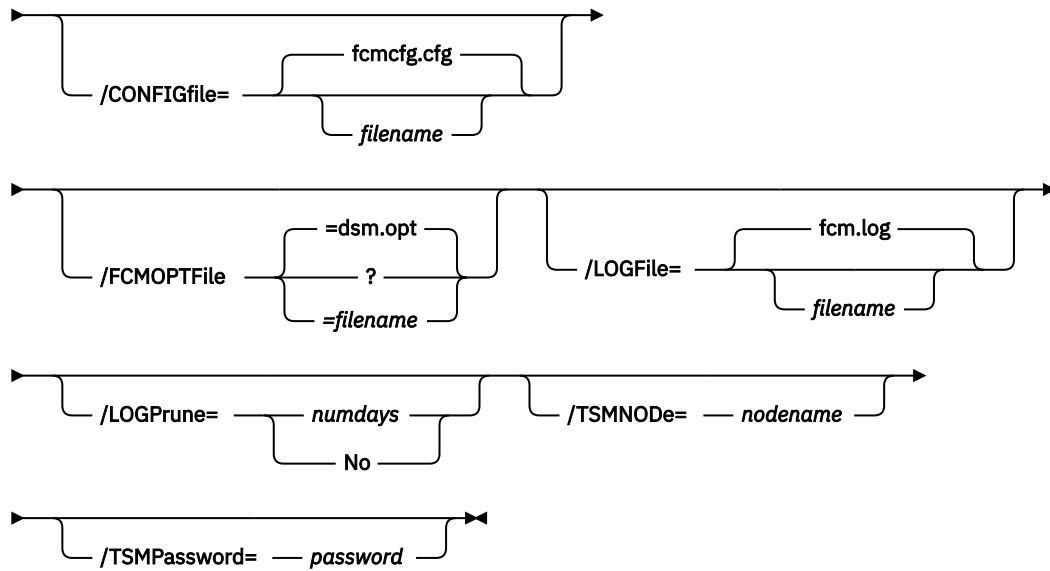
- `policy_name` (required): Specifies the name of the policy that is being created. To create a policy, the policy name must be unique.
- `/DAYSretain`: Specifies the number of days to retain a snapshot (0 - 9999). You can also specify `NOLimit` to represent an unlimited number of days to retain snapshot versions.
- `/VERSretain`: Specifies the number of snapshot versions to retain (1 - 9999). You can also specify `NOLimit` to represent an unlimited number of snapshot versions to retain.

Copy Policy

This command is used to copy an existing policy to a new policy.

FCMCLI command

► FCMCLI — COPY POLICY — *existingPolicyName* — *newPolicyName* →



Parameters:

- *existing_policy_name* (required): Specifies the name of the policy that is being copied.
- *new_policy_name* (required): Specifies the name of the new policy. The policy name must be unique.

Query Policy

This command is used to list the attributes of a policy.

FCMCLI command

► FCMCLI — Query POLICY — *policyName* ➤

For example:

```
c:\Program Files\Tivoli\FlashCopyManager>fcmcli q pol T1

FlashCopy Manager for Windows:
IBM Tivoli Storage FlashCopy Manager
Version 4, Release 1, Level 2.0
(C) Copyright IBM Corporation 2009, 2015. All rights reserved.

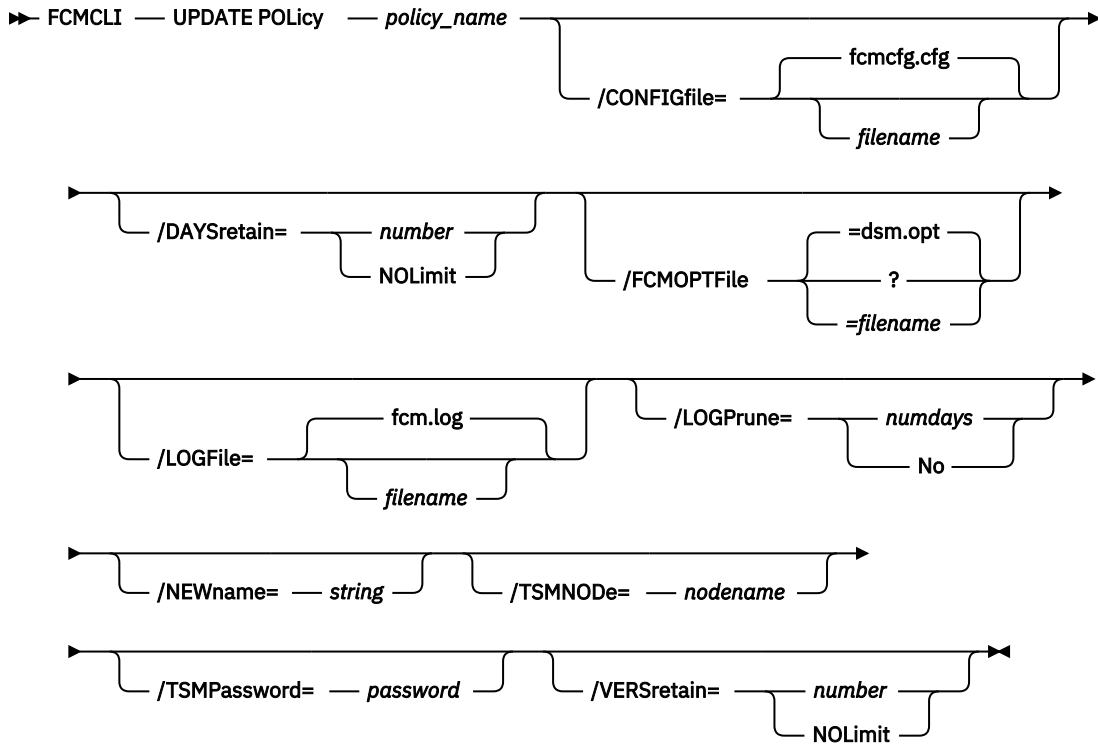
Policy Definitions
=====

Policy Name                : T1
Number of snapshots to keep : No Limit
Number of days to retain snapshot : No Limit
```

Update Policy

This command is used to update or modify an existing policy.

FCMCLI command



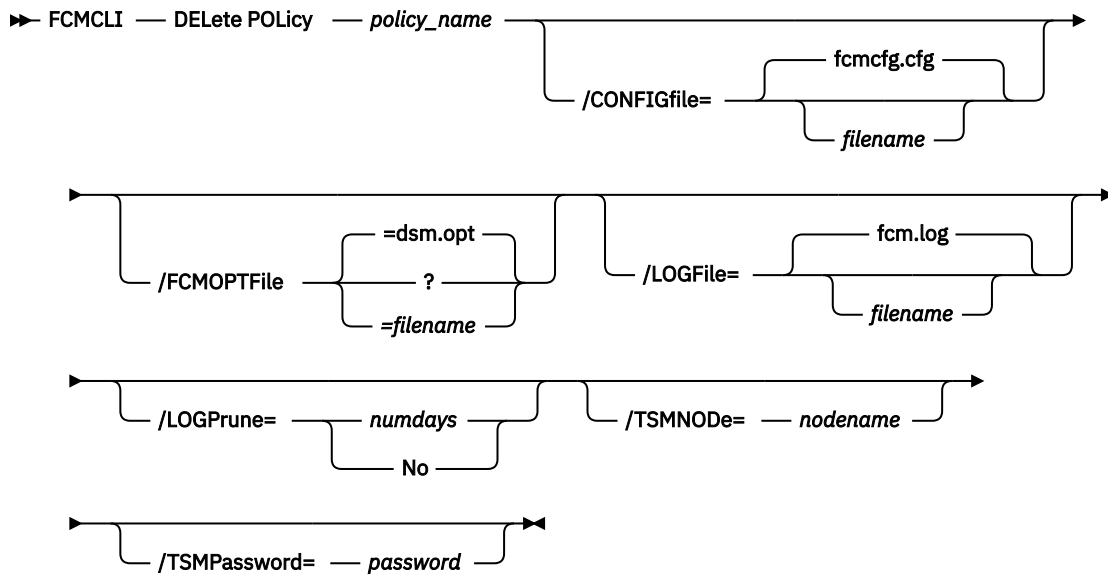
Parameters:

- **NEWname:** Specifies the new name of the policy, if the name is being updated. The policy name must be unique.
- *policy_name* (required): Specifies the name of the policy that is being updated.
- **VERSretain:** Specifies the number of snapshot versions to retain (1 - 9999). You can also specify `NOLimit` to represent an unlimited number of snapshot versions to retain.
- **DAYSretain:** Specifies the number of days to retain a snapshot (0 - 9999). You can also specify `NOLimit` to represent an unlimited number of days to retain snapshot versions.

Delete Policy

This command is used to delete a policy.

FCMCLI command



The required parameter is *policy_name*. The parameter specifies the name of the policy that is being deleted.

IBM Spectrum Protect Snapshot policy examples

These output examples provide a sample of the text, messages, and process status that displays when you use the **create policy** and **delete policy** commands.

In this example, the `fcmlcli create policy FCMPOL1` command creates the *FCMPOL1* policy. The following output is displayed:

```
Policy 'FCMPOL1' was created.  
The operation completed successfully. (rc = 0)  
Completed
```

In this example, the `fcmlcli delete policy FCMPOL1` command deletes the *FCMPOL1* policy. The following output is displayed:

```
Policy 'FCMPOL1' was deleted.  
The operation completed successfully. (rc = 0)  
Completed
```


Mount backup command

Use the **mount backup** command to mount backups that are managed by IBM Spectrum Protect Snapshot or IBM Spectrum Protect.









Mount backup syntax

Use the **mount backup** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command

➤ FCMCLI — MOUNT BACKUP ➔

➤ *component name=mount point* — [, — *component name=mount point* —] ➔



Mount backup positional parameter

The positional parameters immediately follow the **mount backup** command and precede the optional parameters.

The following positional parameters specify the objects to mount:

component name=mount point[,component name=mount point]

component name

Specify the volume or drive name of the component.

mount point

Specify an unused drive letter or absolute path to the directory where the snapshots are going to be displayed as mount point directories. The directory must be empty. If not empty, an error is reported.

The list must contain all non-qualified objects or all qualified objects. The list cannot contain a combination of non-qualified objects and qualified objects. Specify the list by using the following syntax:

```
mount backup object-1[(object-1-id)]= mount-point-1[,object-2[(object-2-id)]
=mount-point-2...]
```

For example:

```
fccli mount backup L:=X:
fccli mount backup g:(2011031112451)=x:
```

The following example is for a non-qualified object `object-1`:

```
fccli mount backup g:=x:
```

The following example is for a qualified object `object-1 (object-1-id)`:

```
fccli mount backup g:(20110815064316)=x:
```

Mount backup optional parameters

Optional parameters follow the **mount backup** command and positional parameters.

/CONFIGfile=configfilename

Use the ***/configfile*** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **mount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the ***/configfile*** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `fcmcfg.xml`.

If the *configfilename* variable includes spaces, enclose the entire ***/configfile*** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

See [“Update config positional parameters” on page 361](#) for descriptions of available configuration parameters.

/FCMOPTfile=filename

Use the *filename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *filename* variable includes spaces, enclose the entire ***/FCMOPTfile*** parameter entry in double quotation marks. For example:

```
/FCMOPTfile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/FROMSERVER=server-name

Use the **/fromserver** parameter to specify the name of the server where the original backup was done. The default is the local server.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `fcm.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is `No`. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

Perform a read-only mount operation.

Yes

Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPMount** parameter in the configuration file.

- If **USESNAPOFASNAPMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
- If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which require IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which require IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate more target volumes on your storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

/PITDate=*date*

Use the **/pitdate** parameter with the **/pittime** parameter to establish a point in time for which you want to mount the latest version of your backups. Backups that were backed up on or before the date and time you specified, and, which were not deleted before the date and time you specified, are processed. Backup versions that you create after this date and time are ignored. Specify the appropriate date in the *date* variable; use the same format that you selected with the DATEFORMAT option in the IBM Spectrum Protect Snapshot options file.

If neither *date* nor *time* is specified, then no date and time are established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup that is selected after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the mount date and time as the current date at the specified *time*.

/PITTime=*time*

Use the **/pittime** parameter with the **/pitdate** option to establish a point in time for which you want to mount the latest version of your backups. Files or images that were backed up on or before the date and time you specify, and that were not deleted before the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify the **/pitdate** parameter. Specify the appropriate time in the *time* variable; use the same format that you selected with the TIMEFORMAT option in the IBM Spectrum Protect Snapshot options file.

If neither *date* nor *time* is specified, then no date and time are established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup that is selected after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the mount date and time as the current date at the specified *time*.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to mount the data.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnodename

Use the *tsmnodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to **PROMPT**. This parameter is not valid when **PASSWORDACCESS** is set to **GENERATE** in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Mount backup examples

These output examples provide a sample of the text, messages, and process status that displays when you use the **mount backup** command.

In this example, the `fcmdi mount backup C:=X:` command mounts volume C:. The following output is displayed:

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a MOUNT BACKUP operation, please wait...

Connecting to TSM Server as node 'STRINGVM1_FS'...
Connecting to Local DSM Agent 'STRINGVM1'...

Backup(s) to be mounted:
C: = X: : VSS : full : 08/04/2016 13:08:50

The operation completed successfully. (rc = 0)
```

In this example, the `fcmdi mount backup D:\mnt\mp1=M:,D:\mnt\mp2=N: /PITDATE=08/07/2014 /PITTIME=08:53:36` command mounts multiple volumes. The following output is displayed:

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a MOUNT BACKUP operation, please wait...

Connecting to FCM Server as node 'TROYVM1_FS'...
Connecting to Local DSM Agent 'TROYVM1'...

Backup(s) to be mounted:

d:\mnt\mp1 = M: : VSS : full : 08/07/2016 08:53:35
d:\mnt\mp2 = N: : VSS : full : 08/07/2016 08:53:36
```

Query component command

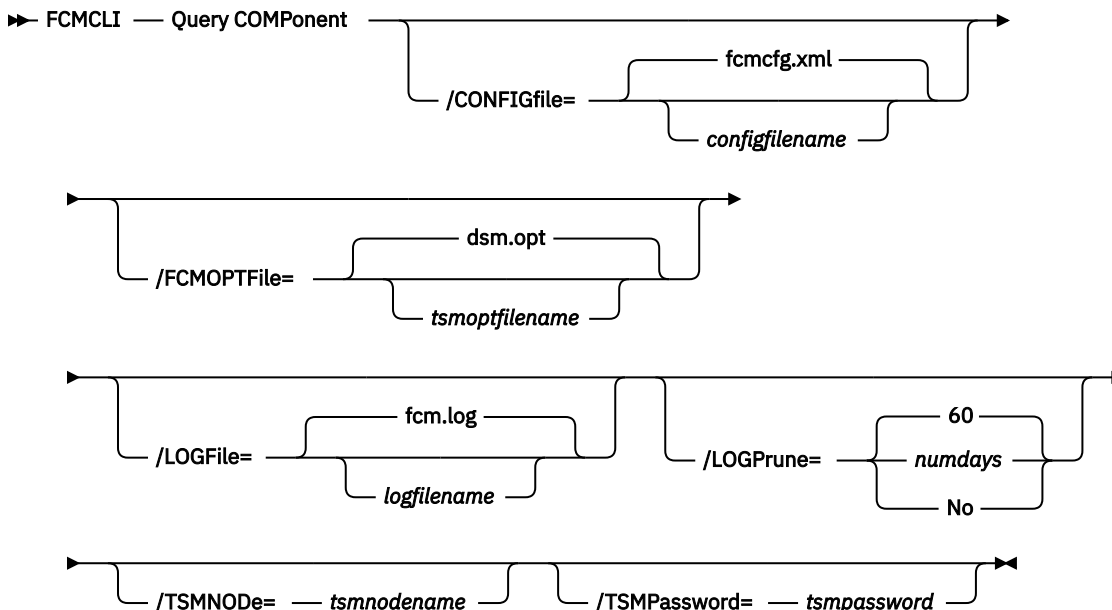
Use the **query component** command to query the VSS components available on the system.

The **query component** command returns a list of the volume and mount points available for backup.

Query component syntax

Use the **query component** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command



Query component optional parameters

Optional parameters follow the **query component** command.

/CONFIGfile=filename

Use the **/configfile** parameter to specify the name (*filename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **query component** operation.

The *filename* variable can include a fully qualified path. If the *filename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *filename* variable is not specified, the default value is `fcmcfg.xml`.

If the *filename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/FCMOPTFile=filename

Use the *filename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *filename* variable includes spaces, enclose the entire **/FCMOPTFile** parameter entry in double quotation marks. For example:

```
/FCMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/LOGFile=filename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *filename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *filename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *filename* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *fcm.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option *No* can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/TSMNode=nodename

Use the *nodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to **PROMPT**. This parameter is not valid when **PASSWORDACCESS** is set to **GENERATE** in the options file.

/TSMPassword=password

Use the *password* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If `PASSWORDACCESS PROMPT` is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Query component examples

Examples of how to use the **query component** command are provided.

To query components that are associated with a configuration file, for example, `customconfig.xml`, enter the following command:

```
fccli query component /configfile=customconfig.xml
```

To query components for a list of the volume and mount points that are available for backup, enter the following command:

```
fccli query component
```

Query config command

Use the **query config** command to display IBM Spectrum Protect Snapshot configuration information.

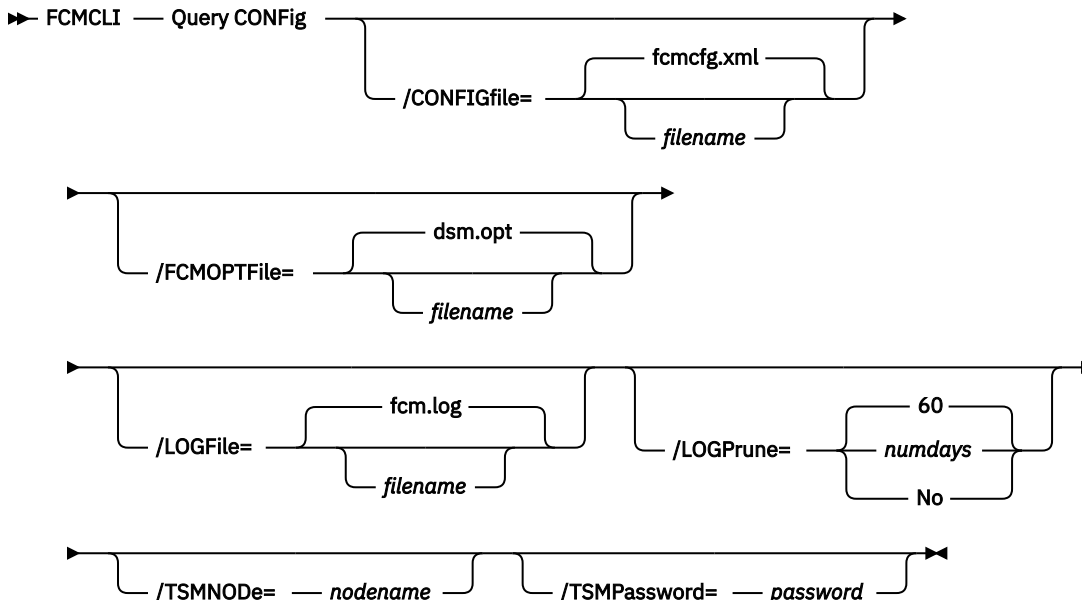
The **query config** command displays the following information:

- The value of each **configuration parameters** parameter
- IBM Spectrum Protect Snapshot connection and configuration information
- IBM Spectrum Protect server connection and configuration information

Query config syntax

Use the **query config** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command



Query config optional parameters

Optional parameters follow the **query config** command.

/CONFIGfile=filename

Use the **/CONFIGfile** parameter to specify the name (*filename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **query config** operation.

The *filename* variable can include a fully qualified path. If the *filename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *filename* variable is not specified, the default value is `fcmcfg.xml`.

If the *filename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/FCMOPTFile=filename

Use the *filename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *filename* variable includes spaces, enclose the entire **/FCMOPTFile** parameter entry in double quotation marks. For example:

```
/FCMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/LOGFile=filename

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *filename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *filename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *filename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, `fcm.log`.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of

the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/TSMNODE=nodename

Use the *nodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMPassword=password

Use the *password* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

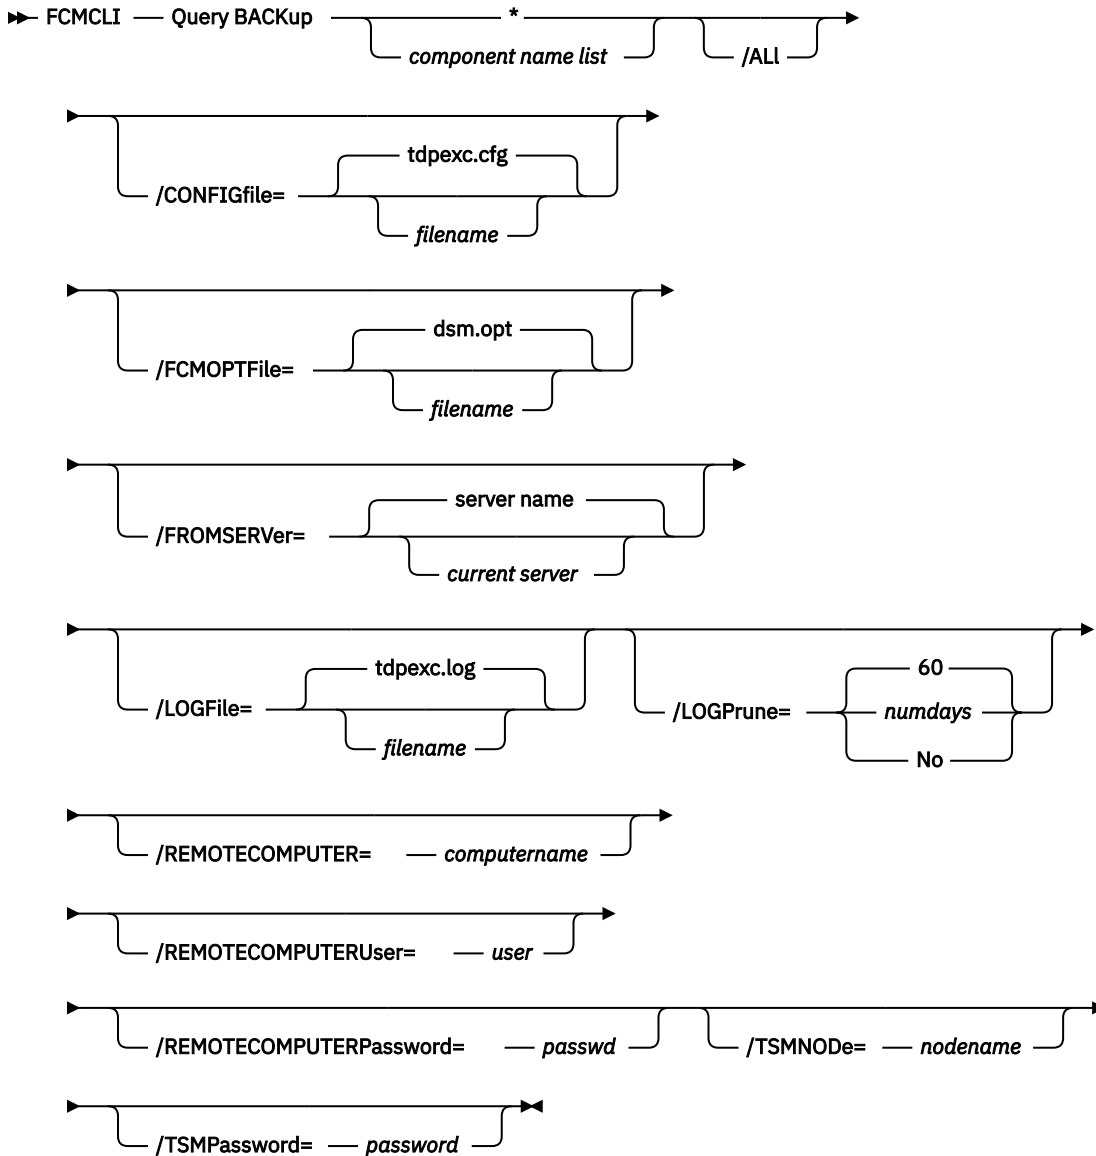
Query backup command

Use the **query backup** command to query a list of the backups that are being managed by IBM Spectrum Protect Snapshot and the IBM Spectrum Protect server.

Query backup syntax

Use the **query backup** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command



Query backup positional parameter

The positional parameter immediately follows the **query backup** command and precedes the optional parameters.

Specify the following positional parameters with the **query backup** command:

component name list | *

component name list

Specify a list of volume or mount points to query.

All backups are queried and shown in the command output. This option is the default value.

Query backup optional parameters

Optional parameters follow the **query backup** command and positional parameter.

/ALL

Use the **/all** parameter to display both active and inactive backup objects. If the **/all** parameter is not specified, only active backup objects are displayed.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **query backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `fcmcfg.xml`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/FCMOPTfile=filename

Use the *filename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *filename* variable includes spaces, enclose the entire **/FCMOPTfile** parameter entry in double quotation marks. For example:

```
/FCMOPTfile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/FROMSERVER=server name

Use the **/fromserver** parameter to specify the name of the server where the original backup was done. The default is the current server.

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `fcm.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify *no*, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/REMOTECOMPUTER=computername

Enter the IP address or host name for the remote system where you want to query the data that is backed up.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=nodename

Use the *nodename* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMPassword=password

Use the *password* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Query backup example

This output example provides a sample of the text, messages, and process status that displays when you use the **query backup** command.

The **fcmcli query backup * /all** command displays information about all active and inactive backups that are managed by IBM Spectrum Protect Snapshot. An example of the output is provided.

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Querying backups, please wait...

Connecting to FCM Server as node 'JUNE_FS'...
Connecting to Local DSM Agent 'JUNE'...

Backups for Volume/Mount Point: 'F:'
=====
Volume/Mount Point      : F:
Volume GUID             : aa3683af-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.13MB
Backup Date/Time        : 03/31/2016 07:35:11
Backup State             : Active
Management Class        : DEFAULT
Mounted as              :
Object Name              : 20140331073511
Instant Restore Supported : No

Volume/Mount Point      : F:
Volume GUID             : aa3683af-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.13MB
Backup Date/Time        : 03/30/2016 13:50:44
Backup State             : Inactive
Management Class        : DEFAULT
Mounted as              :
Object Name              : 20140330135044
Instant Restore Supported : No

Backups for Volume/Mount Point: 'O:'
=====
Volume/Mount Point      : O:
Volume GUID             : aa3683b2-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.12MB
Backup Date/Time        : 03/31/2016 07:35:50
Backup State             : Active
Management Class        : DEFAULT
Mounted as              :
Object Name              : 20140331073550
Instant Restore Supported : No

Volume/Mount Point      : O:
Volume GUID             : aa3683b2-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.12MB
Backup Date/Time        : 03/31/2016 07:24:44
Backup State             : Inactive
Management Class        : DEFAULT
Mounted as              :
Object Name              : 20140331072444
Instant Restore Supported : No
```

The **fcmcli query backup** command displays information about backups that are managed by IBM Spectrum Protect Snapshot. An example of the output is provided.


```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.
```

Querying backups, please wait...

```
Connecting to FCM Server as node 'JUNE_FS'...
Connecting to Local DSM Agent 'JUNE'...
```

Backups for Volume/Mount Point: 'F:'

```
=====
Volume/Mount Point      : F:
Volume GUID             : aa3683af-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.13MB
Backup Date/Time        : 03/31/2016 07:35:11
Backup State            : Active
Management Class        : DEFAULT
Mounted as              :
Object Name             : 20140331073511
Instant Restore Supported : No
```

Backups for Volume/Mount Point: '0:'

```
=====
Volume/Mount Point      : 0:
Volume GUID             : aa3683b2-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.12MB
Backup Date/Time        : 03/31/2016 07:35:50
Backup State            : Active
Management Class        : DEFAULT
Mounted as              :
Object Name             : 20140331073550
Instant Restore Supported : No
```

The **fcmlcli query backup** command displays information about backups that are managed by IBM Spectrum Protect Snapshot. An example of the output is provided.

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.
```

Querying backups, please wait...

```
Connecting to FCM Server as node 'JUNE_FS'...
Connecting to Local DSM Agent 'JUNE'...
```

Backups for Volume/Mount Point: 'F:'

```
=====
Volume/Mount Point      : F:
Volume GUID             : aa3683af-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.13MB
Backup Date/Time        : 03/31/2016 07:35:11
Backup State            : Active
Management Class        : DEFAULT
Mounted as              :
Object Name             : 20140331073511
Instant Restore Supported : No
```

Backups for Volume/Mount Point: '0:'

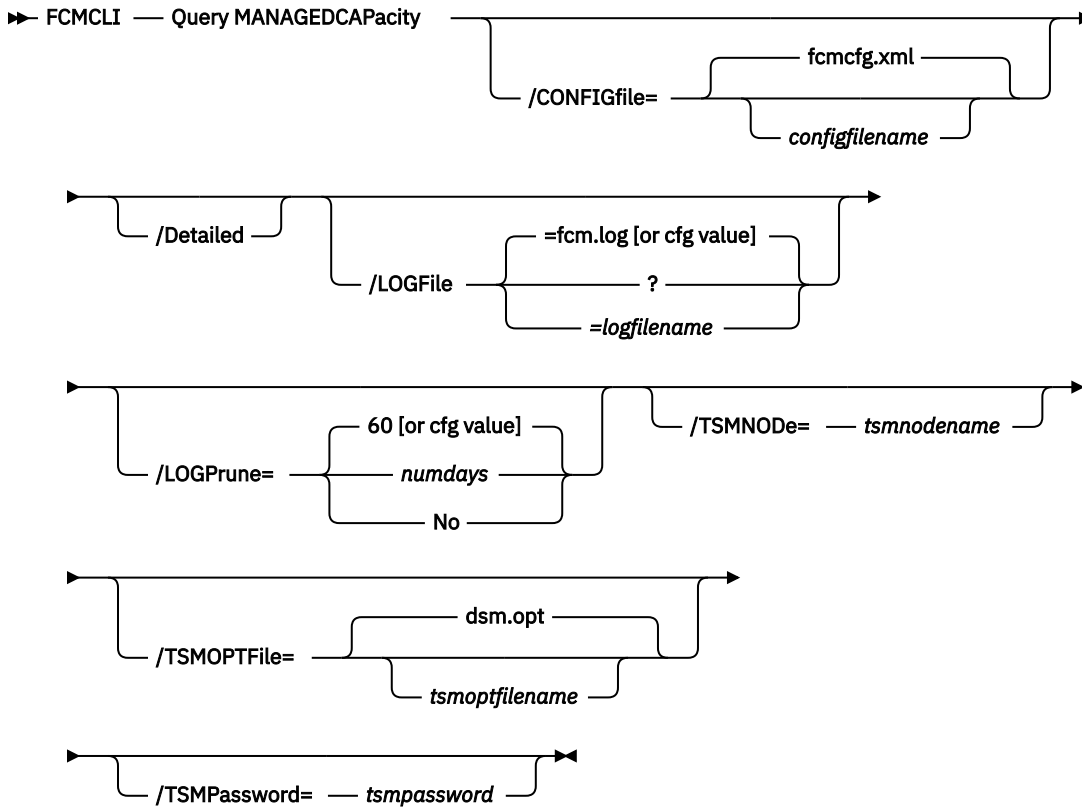
```
=====
Volume/Mount Point      : 0:
Volume GUID             : aa3683b2-4bdc-11de-b146-001a6499a400
Server                  : JUNE
Volume Occupancy        : 10.12MB
Backup Date/Time        : 03/31/2016 07:35:50
Backup State            : Active
Management Class        : DEFAULT
Mounted as              :
Object Name             : 20140331073550
Instant Restore Supported : No
```

Query managedcapacity command

When you plan for storage, if you want to determine the amount of managed capacity in use, use the **query managedcapacity** command.

The **query managedcapacity** command displays capacity that is related information about the volumes that are represented in local inventory that is managed by IBM Spectrum Protect Snapshot. You can run this command on all Windows operating systems that are supported by IBM Spectrum Protect Snapshot.

FCMCLI command



Parameters

/CONFIGfile=*configfilename*

Use the **/CONFIGfile** parameter to specify the name (*configfilename*) of the configuration file that contains the values to use for a **query managedcapacity** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/CONFIGfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is *fcmcfg.xml*.

If the *configfilename* variable includes spaces, enclose the entire **/CONFIGfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/Detailed

Results in a detailed listing of snapped volumes. If this option is not specified, only the total capacity is displayed.

/LOGFile=*logfile*

Use the **/LOGFile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfile* variable includes spaces, enclose the entire **/LOGFile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/LOGFile** parameter is not specified, log records are written to the default log file, *fcm.log*.

The **/LOGFile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to **PROMPT**. This parameter is not valid when **PASSWORDACCESS** is set to **GENERATE** in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

For examples of how to use the **query managedcapacity** command, use the following list:

- To display the total amount of managed capacity in use in the local inventory, enter the **fcmlcli query managedcapacity** command.

If there are local backups, the following code sample can be used as a reference:

```
c:\Program Files\Tivoli\FlashCopyManager>fcmlcli query managedcapacity

IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a QUERY MANAGEDCAPACITY operation, please wait...

Total Managed Capacity : 84.26 GB (90,476,371,968 bytes)
```

If there are no local backups, the following code sample can be used as a reference:

```
c:\Program Files\Tivoli\FlashCopyManager>fcmlcli query managedcapacity

IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a QUERY MANAGEDCAPACITY operation, please wait...

Total Managed Capacity : 0
```

- To display a detailed listing of total amount of managed capacity and the snapped volumes in use, enter the **fcmlcli query managedcapacity /detailed** command.

If there are local backups, the following code sample can be used as a reference:

```
c:\Program Files\Tivoli\FlashCopyManager>fcmlcli query managedcapacity /detail

IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a QUERY MANAGEDCAPACITY operation, please wait...

Total Managed Capacity : 84.26 GB (90,476,371,968 bytes)

Volume          : C:
Managed Capacity : 68.27 GB (73,299,652,608 bytes)
```

```
Volume      : c:\mp
Managed Capacity : 16.00 GB (17,176,719,360 bytes)
```

If there are no local backups, the following code sample can be used as a reference:

```
c:\Program Files\Tivoli\FlashCopyManager>fcmcli query managedcapacity /detail

IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a QUERY MANAGEDCAPACITY operation, please wait...

Total Managed Capacity : 0
```

Restore command

Use the **restore** command to restore an IBM Spectrum Protect Snapshot backup.

You must have local registry rights to run an IBM Spectrum Protect Snapshot for Exchange Server restore.

VSS operations require special considerations that must be reviewed before you attempt a VSS Restore. See these two sections for important guidelines:

- [“VSS restore characteristics” on page 3](#)
- [“VSS backups that are restored to alternate databases” on page 18](#)

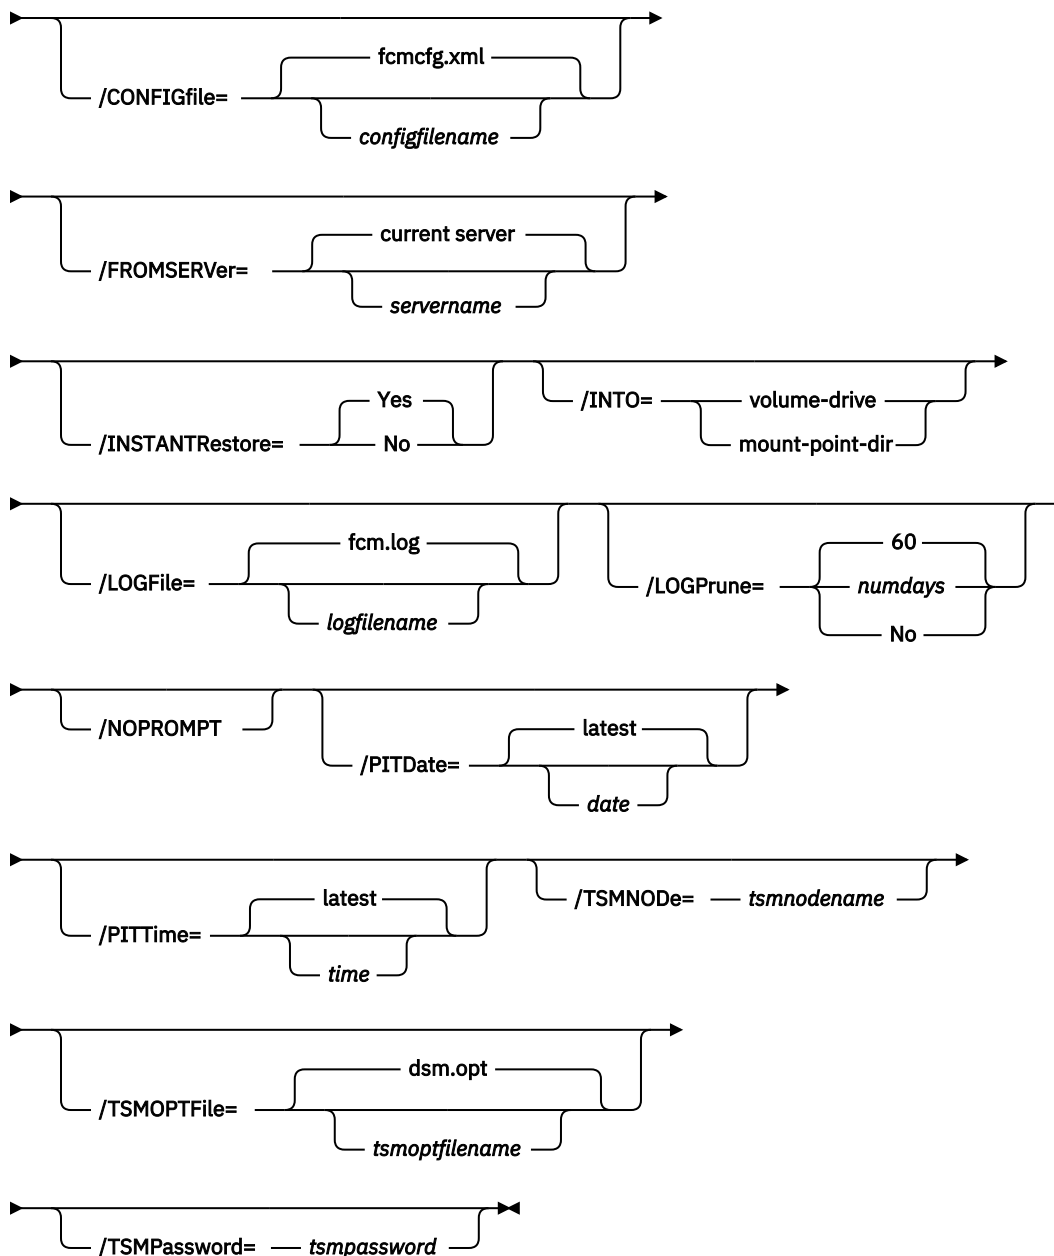
The GUI provides an easy-to-use, flexible interface to help you run a restore operation. The interface presents information in a way that allows multiple selection and, in some cases, automatic operation.

Restore syntax

Use the **restore** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command

➤ FCMCLI — RESTore — *component name list* ➔



Restore positional parameter

The positional parameter immediately follows the **restore** command and precedes the optional parameters.

Specify the following positional parameter with the **restore** command:

component name list

Specify a list of volume or mount points to restore. The list must contain all non-qualified objects or all qualified objects. The list cannot contain a combination of non-qualified objects and qualified objects.

Specify the component name list by using the following syntax:

```
comp-1[(object-1-id)][,comp-2[(object-2-id)]...]
```

where *comp-n* is the component to restore, and *obj-id-n* is the object ID of the specific backup to restore. The object ID can be obtained through the **query backup** command.

For example:

```
fcmcli restore g:(20110311124516),h:(20110211034512),r:(20101114164310)
```

Restore optional parameters

Optional parameters follow the **restore** command and positional parameters.

/CONFIGfile=configfilename

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for a **restore** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `fcmcfg.xml`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

See [“Update config positional parameters” on page 361](#) for descriptions of available configuration parameters.

/FROMSERVER=server-name

Use the **/fromserver** parameter to specify the name of the server where the original backup was done. The default is the local server.

/INSTANTRestore=Yes | No

Use the **/instantrestore** parameter to specify whether to use volume level snapshot or file level copy to restore a VSS backup that are on local shadow volumes. An IBM Systems Storage SAN Volume Controller, DS8000, the XIV system, and IBM Storwize V7000 storage subsystem is required to perform VSS instant restores.

You can specify:

Yes

Use volume level snapshot restore for a VSS backup on local shadow volumes if the backup exists on volumes that support it. This option is the default.

No

Use file level copy to restore the files from a VSS backup on local shadow volumes. Bypassing volume-level copy means that Exchange log files and the checkpoint file are the only data overwritten on the source volumes.

When you complete VSS instant restores with DS8000, make sure that any previous background copies (that involve the volumes that are being restored) are completed before you initiate the VSS instant restore. The **/instantrestore** parameter is ignored and VSS instant restore capabilities are automatically disabled when performing any type of VSS restore into operation. You cannot run VSS instant restore of differential and incremental backups.

/INTO=volume-drive | mount-point-dir

Use the **/into** parameter to restore the backup that is stored on IBM Spectrum Protect server to an alternate destination.

You can specify either *volume-drive* or *mount-point-dir*. The *volume-drive* or *mount-point-dir* location that you specify must be present on the server; the location is not dynamically created.

You can issue the **/into** parameter for one restore operation per command. You cannot use multiple restore specifications with the **/into** parameter.

The following sample provides an example of how to use the parameter:

```
FCMCLI RESTORE M: /INTO=P:
```

/LOGFile=logfilename

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfilename* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `fcm.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option `No` can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify `no`, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/NOPROMPT

When the **restore** command is issued, you are prompted to confirm whether to overwrite the volumes you specified for restore. Use the **/noprompt** parameter to bypass this prompt and proceed with the restore operation.

/PITDate=date

Use the ***/pitdate*** parameter with the ***/pittime*** parameter to establish a point in time for which you want to restore the latest version of your backups. Backups that were backed up on or before the date and time you specified, and, which were not deleted before the date and time you specified, are processed. Backup versions that you create after this date and time are ignored. Specify the appropriate date in the *date* variable; use the same format that you selected with the DATEFORMAT option in the IBM Spectrum Protect Snapshot options file.

If neither *date* nor *time* is specified, then no date and time are established. By default the backup is restored from the most recent available backup.

If either *date* or *time* is specified, then the backup is restored from the earliest backup that is selected after the established restore date and time. If no backup after the established date and time is found, by default the backup is restored from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the restore period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the restore date and time as the current date at the specified *time*.

/PITTime=time

Use the ***/pittime*** parameter with the ***/pitdate*** option to establish a point in time for which you want to restore the latest version of your backups. Files or images that were backed up on or before the date and time you specify, and that were not deleted before the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify the ***/pitdate*** parameter. Specify the appropriate time in the *time* variable; use the same format that you selected with the TIMEFORMAT option in the IBM Spectrum Protect Snapshot options file.

If neither *date* nor *time* is specified, then no date and time are established. By default the backup is restored from the most recent available backup.

If either *date* or *time* is specified, then the backup is restored from the earliest backup that is selected after the established restore date and time. If no backup after the established date and time is found, by default the backup is restored from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the restore period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the restore date and time as the current date at the specified *time*.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is `dsm.opt`.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified `PASSWORDACCESS GENERATE` in the IBM Spectrum Protect Snapshot options file (`dsm.opt`), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when `PASSWORDACCESS GENERATE` is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If `PASSWORDACCESS PROMPT` is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Restore examples

These output examples provide a sample of the text, messages, and process status that displays when you use the **restore** command.

In this example, the `fcmdi restore K:,L: /INSTANTRestore=No` command restores volumes K: and L:. The following output is displayed:

```
IBM Spectrum Protect Snapshot for Windows:  
IBM Spectrum Protect Snapshot  
Version 8, Release 1, Level 0  
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.
```

```
You have selected a full filesystem RESTORE operation. Performing this restore  
will overwrite the volumes that you have specified for restore.
```

```
Do you want to continue with the RESTORE operation? (Yes (Y)/No (N)) y
```

```
Preparing for a RESTORE operation, please wait...
```

```
Starting restore of volume...
```

```
Beginning VSS restore of 'K:', 'L:'. This operation could take a while, please wait...
```

```
Restoring 'K:', 'L:' via file-level copy from snapshot(s). This process may take  
some time. Please wait.
```

```
VSS Restore operation completed with rc = 0.
```

```
Elapsed Processing Time: 385.23 seconds
```

In this example, the `fcmdi restore D:\mnt\mp1,D:\mnt\mp2 /PITDATE=10/07/2014 /PITTIME=08:53:36` command restores mount points D:\mnt\mp1 and ,D:\mnt\mp2. The following output is displayed:

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.
```

You have selected a full filesystem RESTORE operation. Performing this restore will overwrite the volumes that you have specified for restore.

Do you want to continue with the RESTORE operation? (Yes (Y)/No (N)) y

Preparing for a RESTORE operation, please wait...

Starting restore of volume...

Beginning VSS restore of 'd:\mnt\mp1', 'd:\mnt\mp2'. This operation could take a while, please wait...

Restoring 'd:\mnt\mp1', 'd:\mnt\mp2' via volume-level copy from snapshot(s). This process may take some time. Please wait.

VSS Restore operation completed with rc = 0.

Elapsed Processing Time: 162.23 seconds

In this example, the `fcmdi restore K:,L: /FROMSERVER=troyvm1` command restores volumes K: and L: from server troyvm1. The following output is displayed:

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.
```

You have selected a full filesystem RESTORE operation. Performing this restore will overwrite the volumes that you have specified for restore.

Do you want to continue with the RESTORE operation? (Yes (Y)/No (N)) y

Preparing for a RESTORE operation, please wait...

Starting restore of volume...

Beginning VSS restore of 'K:', 'L:'. This operation could take a while, please wait...

Restoring 'K:', 'L:' via volume-level copy from snapshot(s). This process may take some time. Please wait.

VSS Restore operation completed with rc = 0.

Elapsed Processing Time: 161.57 seconds

Unmount backup command

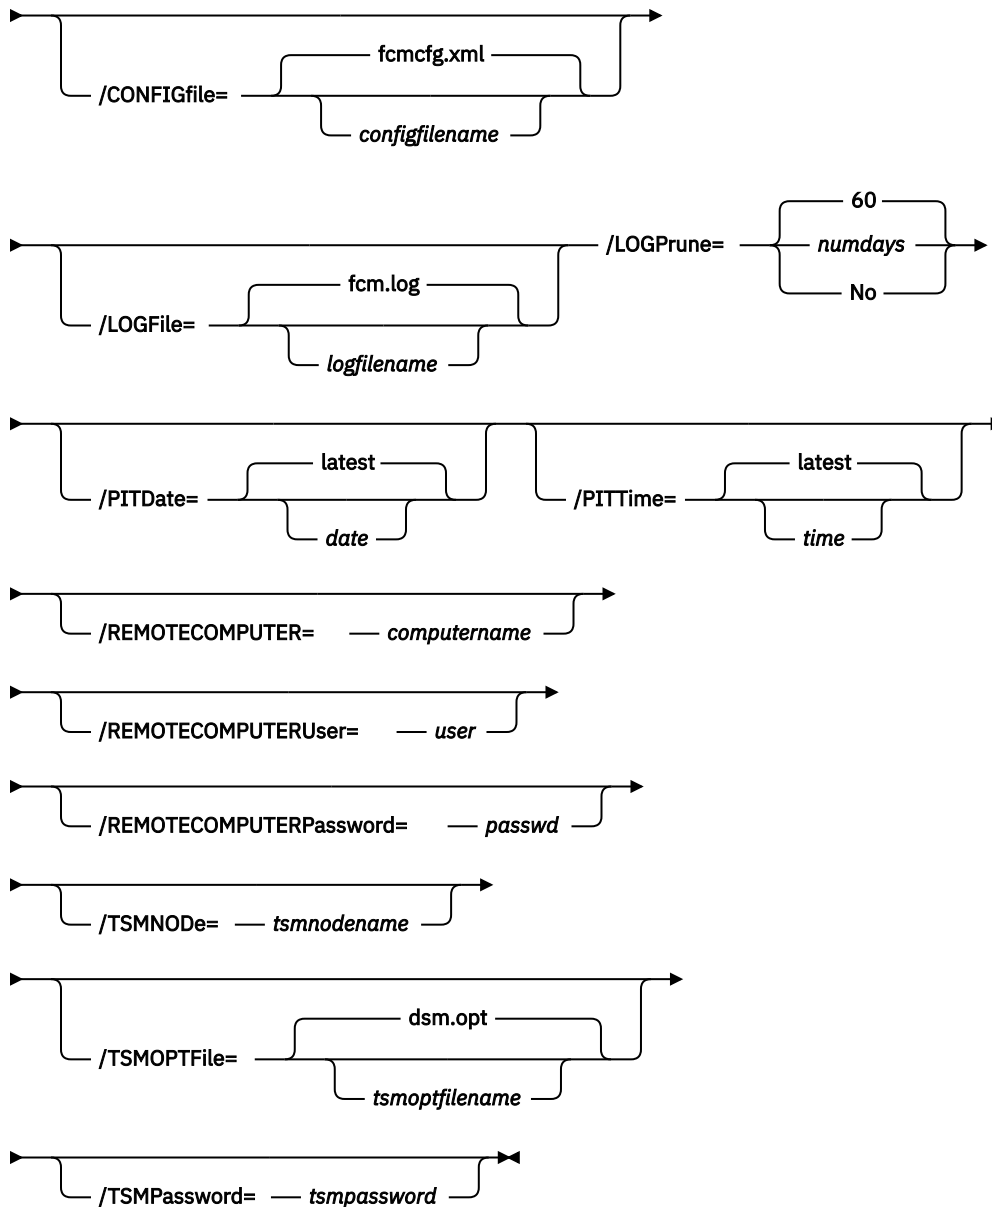
Use the **unmount backup** command to unmount backups that were previously mounted, and are managed by IBM Spectrum Protect Snapshot or IBM Spectrum Protect.

Unmount backup syntax

Use the **unmount backup** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command

► FCMCLI — UNMOUNT BACKUP — *volume name list* ►



Unmount backup positional parameter

The positional parameter immediately follows the **unmount backup** command and precedes the optional parameters.

volume name list

Use this parameter to specify a drive letter (for example, a:) or list of mount point directories to unmount. The *volume name list* parameter is required.

To specify more than one name, separate them by commas.

Unmount backup optional parameters

Optional parameters follow the **unmount backup** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for an **unmount backup** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is `fcmcfg.xml`.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/LOGFile=*logfile*

Use the **/logfile** parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The *logfile* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfile* variable can include a fully qualified path. However, if no path is specified, the log file is written to the IBM Spectrum Protect Snapshot installation directory.

If the *logfile* variable includes spaces, enclose the entire **/logfile** parameter entry in double quotation marks. For example:

```
/LOGFile="c:\Program Files\myfcm.log"
```

If the **/logfile** parameter is not specified, log records are written to the default log file, `fcm.log`.

The **/logfile** parameter cannot be turned off, logging always occurs.

/LOGPrune=*numdays***|No**

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, a certain number of days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.
- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is 60.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, 60, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of

the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

/PITDate=date

Use the **/pitdate** parameter with the **/pittime** parameter to establish a point in time for which you want to mount the latest version of your backups. Backups that were backed up on or before the date and time you specified, and, which were not deleted before the date and time you specified, are processed. Backup versions that you create after this date and time are ignored. Specify the appropriate date in the *date* variable; use the same format that you selected with the DATEFORMAT option in the IBM Spectrum Protect Snapshot options file.

If neither *date* nor *time* is specified, then no date and time are established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup that is selected after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the mount date and time as the current date at the specified *time*.

/PITTime=time

Use the **/pittime** parameter with the **/pitdate** option to establish a point in time for which you want to mount the latest version of your backups. Files or images that were backed up on or before the date and time you specify, and that were not deleted before the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify the **/pitdate** parameter. Specify the appropriate time in the *time* variable; use the same format that you selected with the TIMEFORMAT option in the IBM Spectrum Protect Snapshot options file.

If neither *date* nor *time* is specified, then no date and time are established. By default the backup is mounted from the most recent available backup.

If either *date* or *time* is specified, then the backup is mounted from the earliest backup that is selected after the established mount date and time. If no backup after the established date and time is found, by default the backup is mounted from the most recent available backup.

Notes:

- If you specify both *date* and *time*, this selection establishes the mount backup period.
- If you specify *date* and you do not specify *time*, *time* defaults to a value of 23:59:59. This selection establishes the *date* at the specified date.
- If you specify *time* without *date*, then *date* defaults to the current date. This selection establishes the mount date and time as the current date at the specified *time*.

/REMOTECOMPUTER=computername

Enter the computer name or IP address of the remote system where the backup was created.

/REMOTECOMPUTERUser=user

Enter the user name that is used to log on to the server specified with the **REMOTECOMPUTER** parameter. If a domain is required to log on with the domain account, enter *Domain\User*. To log on to the local account, the domain is not required. There is no default value.

/REMOTECOMPUTERPassword=passwd

Enter the password for the user name that is specified with the **REMOTECOMPUTERUser** parameter. There is no default value.

/TSMNODE=tsmnode

Use the *tsmnode* variable to refer to the IBM Spectrum Protect node name that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

You can store the node name in the IBM Spectrum Protect options file (*dsm.opt*). This parameter overrides the value in the IBM Spectrum Protect options file if **PASSWORDACCESS** is set to PROMPT. This parameter is not valid when **PASSWORDACCESS** is set to GENERATE in the options file.

/TSMOPTFile=tsmoptfilename

Use the *tsmoptfilename* variable to identify the IBM Spectrum Protect options file.

The file name can include a fully qualified path name. If no path is specified, the directory where IBM Spectrum Protect Snapshot is installed is searched.

If the *tsmoptfilename* variable includes spaces, enclose the entire **/tsmoptfile** parameter entry in double quotation marks. For example:

```
/TSMOPTFile="c:\Program Files\file.opt"
```

The default is *dsm.opt*.

/TSMPassword=tsmpassword

Use the *tsmpassword* variable to refer to the IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server.

If you specified **PASSWORDACCESS GENERATE** in the IBM Spectrum Protect Snapshot options file (*dsm.opt*), supplying the password is not necessary here because the one that is stored in the registry is used. However, to store the password in the registry, you must specify the IBM Spectrum Protect password the first time IBM Spectrum Protect Snapshot connects to the IBM Spectrum Protect server.

If you do specify a password with this parameter when **PASSWORDACCESS GENERATE** is in effect, the command-line value is ignored unless the password for this node is not yet stored in the registry. In that case, the specified password is stored in the registry and used when you run this command.

If **PASSWORDACCESS PROMPT** is in effect, and you do not specify a password value on the command line, then you are prompted for a password.

The IBM Spectrum Protect password that IBM Spectrum Protect Snapshot uses to log on to the IBM Spectrum Protect server can be up to 63 characters in length.

Unmount backup example

This output example provides a sample of the text, messages, and process status that displays when you use the **unmount backup** command.

In this example, the `fcmdi unmount backup M: ,N:` command unmounts mount points M: and N: The following output is displayed:

```
IBM Spectrum Protect Snapshot for Windows:
IBM Spectrum Protect Snapshot
Version 8, Release 1, Level 0
(C) Copyright IBM Corporation 2009, 2016. All rights reserved.

Preparing for a UNMOUNT BACKUP operation, please wait...

Connecting to FCM Server as node 'TROYVM1_FS'...
Connecting to Local DSM Agent 'TROYVM1'...

Backup(s) to be unmounted:
M:
N:

The operation completed successfully. (rc = 0)
```

Update config command

Use the **update config** command to set the IBM Spectrum Protect Snapshot or IBM Spectrum Protect configuration parameters in a configuration file.

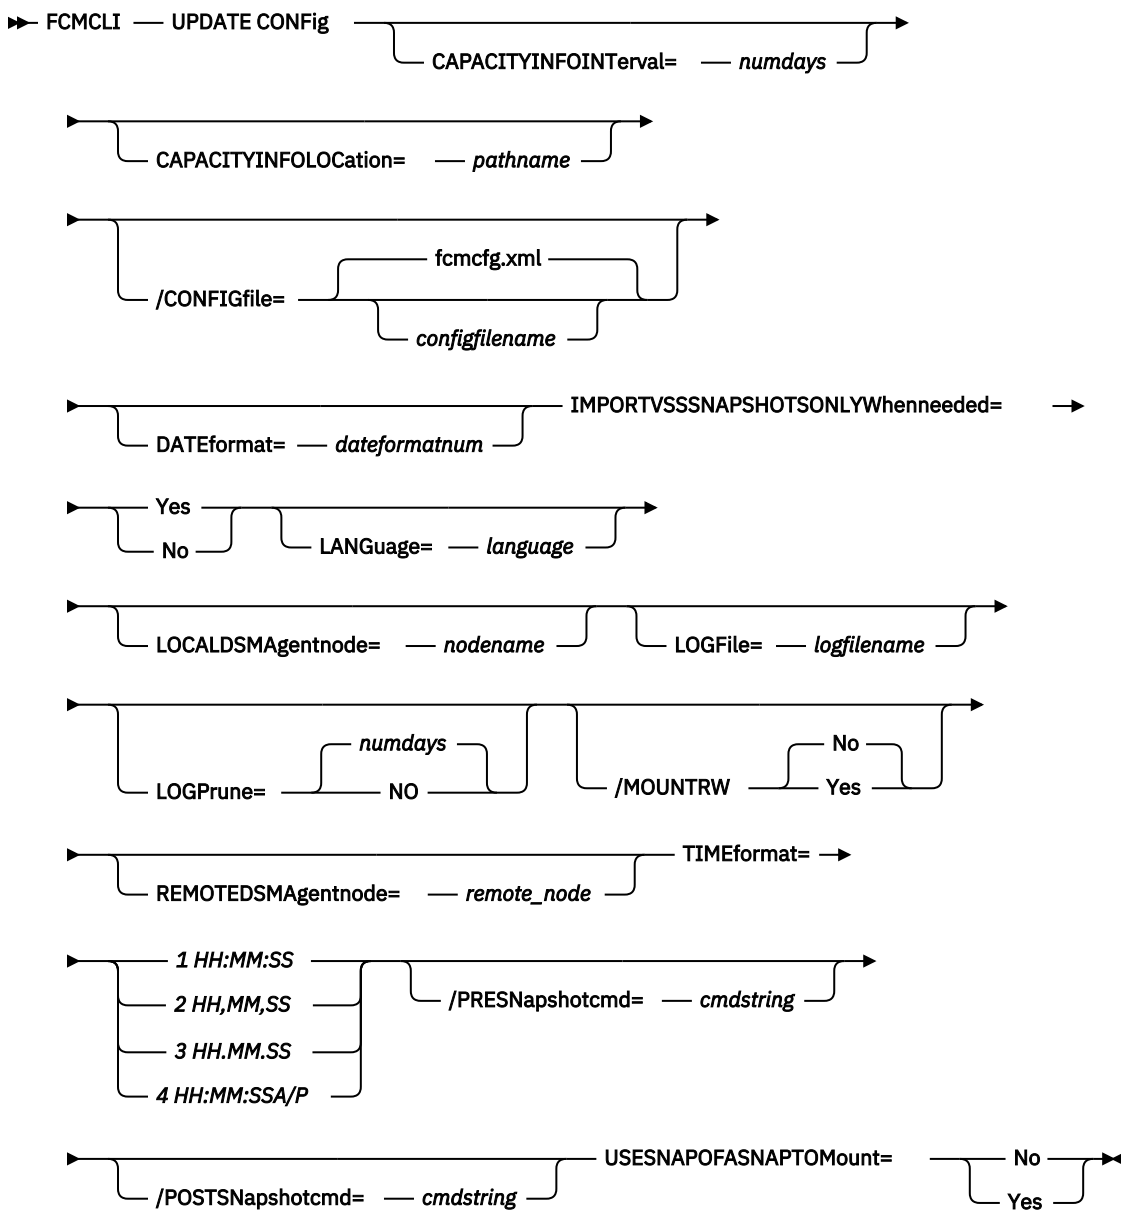
The values that you set are saved in an IBM Spectrum Protect Snapshot configuration file. The default configuration file is `fcmcfg.xml`. Configuration values can also be set in the **Properties** window in Microsoft Management Console (MMC).

For command invocations other than this command, the value of a configuration parameter that is specified in a command overrides the value of the configuration parameter that is specified in the IBM Spectrum Protect Snapshot configuration file. If, when you use this command, you do not override a value for the configuration file parameter, the values in the default configuration file are used.

Update config syntax

Use the **update config** command syntax diagrams as a reference to view available options and truncation requirements.

FCMCLI command



Update config positional parameters

Positional parameters immediately follow the **update config** command and precede the optional parameters.

The following positional parameters specify the values in the IBM Spectrum Protect Snapshot configuration file. You can set only one value for each **update config** command run:

CAPACITYINFOInterval=numdays

Use the **CAPACITYINFOInterval** positional parameter to specify how often you want the capacity metrics report to be generated.

The report, in an XML file format, is generated automatically at the end of a backup operation. The valid value range is 1 - 365 and the default value is 7 days, which means the report is generated once every 7 days.

CAPACITYINFOLOcation=pathname

Use the **CAPACITYINFOLOcation** positional parameter to specify the location where you want the capacity metrics report to be created. If you do not specify a location, the report is not generated.

DATEformat=dateformatnum

Use the **DATEformat** positional parameter to select the format you want to use to display dates.

The *dateformatnum* variable displays the date in one of the following formats. Select the format number that corresponds to the format you want to use.

- 1** MM/DD/YYYY. This format is the default.
- 2** DD-MM-YYYY
- 3** YYYY-MM-DD
- 4** DD.MM.YYYY
- 5** YYYY.MM.DD
- 6** YYYY/MM/DD
- 7** DD/MM/YYYY

Changes to the value of the **dateformat** parameter can result in an undesired pruning of the IBM Spectrum Protect Snapshot log file (`fcml.log` by default). You can avoid losing existing log file data by doing one of the following actions:

- After you change the value of the **dateformat** parameter, make a copy of the existing log file before you run IBM Spectrum Protect Snapshot.
- Specify a new log file with the **/logfile** parameter.

IMPORTVSSSNAPSHOTSONLYWhenneeded

Use the **/IMPORTVSSSNAPSHOTSONLYWhenneeded** parameter to specify whether IBM Spectrum Protect Snapshot automatically imports VSS snapshots to the Windows system where the snapshots are created.

Specify one of the following values:

Yes

Import VSS snapshots to the Windows system where the snapshots are created. The option is the default. During backup processing, transportable snapshots are automatically created and imported to storage systems when the snapshots are required. This option is the default.

Restriction: For instant restore processing on third-party storage systems, you must specify the Yes option to enable the storage system to create transportable snapshots during backups.

No

Do not create transportable VSS snapshots during backup processing, and do not automatically import the snapshot to storage systems after the backup is completed.

LANGUage=language

Specify the three-character code of the language you want to use to display messages:

CHS

Simplified Chinese

CHT

Traditional Chinese

DEU

Standard German

ENU

American English (This option is the default.)

ESP

Standard Spanish

FRA

Standard French

ITA

Standard Italian

JPN

Japanese

KOR

Korean

PTB

Brazilian Portuguese

LOCALDSMAgentnode=nodename

Specify the node name of the local system that runs the VSS backups. This positional parameter must be specified for VSS operations to be completed.

LOGFile=logfilename

Use the **LOGFile** positional parameter to specify the name of the activity log file that is generated by IBM Spectrum Protect Snapshot. The IBM Spectrum Protect Snapshot activity log records significant events, such as completed commands and error messages.

The *logfilename* variable identifies the name of the activity log file.

If the specified log file does not exist, a new log file is created. If the specified log file exists, new log entries are appended to the file. The *logfilename* variable can include a fully qualified path. However, if no path is specified, the log file is assigned to the IBM Spectrum Protect Snapshot installation directory.

/LOGPrune=numdays|No

When you prune log data, you can discard some of the generated logs according to detailed filtering criteria that you set. Depending on the option that you set for the **/LOGPrune** parameter, some days of data are saved. By default, 60 days of log entries are saved. The option **No** can be entered to disable log pruning.

Regardless of the option that you set for this parameter, you can explicitly request log pruning at any time.

Considerations:

- For *numdays*, the range is 0 to 9999. A value of 0 deletes all entries in the activity log file except for the current command entries.
- If you specify **no**, the log file is not pruned.

- If you do not specify **/LOGPrune**, the default value is that specified by the **logprune** configurable option in the configuration file. The default value is *60*.
- If you specify **/LOGPrune**, its value is used instead of the value that is stored in the configuration file. Specifying this parameter does not change the value in the configuration file.
- You can specify **/LOGPrune** without specifying *numdays* or *no*; in this case, the default value, *60*, is used.
- Changes to the value of the **TIMEformat** or **DATEformat** parameter can result in an unwanted pruning of the log file. If you are running a command that might prune the log file, and the value of the **TIMEformat** or **DATEformat** parameter is changed, complete one of the following to prevent unintentional pruning of the log file:
 - Make a copy of the existing log file.
 - Specify a new log file with the **/LOGFile** parameter or **logfile** setting.

NUMberformat=*fmtnum*

Use the **NUMberformat** positional parameter to specify the format you want to use to display numbers.

The *fmtnum* variable displays numbers by using one of the following formats. Select the format number that corresponds to the format you want to use.

- 1**
n,nnn.dd. This format is the default.
- 2**
n,nnn,dd.
- 3**
n nnn,dd
- 4**
n nnn.dd
- 5**
n.nnn,dd
- 6**
n'nnn,dd

REMOTEDSMAgentnode=*remote_node*

Specifies the remote client node that runs the VSS offloaded backups on a remote computer.

TIMEformat=*formatnumber*

Use the **TIMEformat** positional parameter to specify the format in which you want to display the system time.

The *formatnumber* variable displays time in one of the following formats. Select the format number that corresponds to the format you want to use.

- 1**
HH:MM:SS This is the default.
- 2**
HH,MM,SS
- 3**
HH.MM.SS
- 4**
HH:MM:SSA/P

Update config optional parameters

Optional parameters follow the **update config** command and positional parameters.

/CONFIGfile=*configfilename*

Use the **/configfile** parameter to specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use for an **update config** operation.

The *configfilename* variable can include a fully qualified path. If the *configfilename* variable does not include a path, the IBM Spectrum Protect Snapshot installation directory is used. If the **/configfile** parameter is not specified, or if the *configfilename* variable is not specified, the default value is *fcmcfg.xml*.

If the *configfilename* variable includes spaces, enclose the entire **/configfile** parameter entry in double quotation marks. For example:

```
/CONFIGfile="c:\Program Files\fcmcfg.xml"
```

/MOUNTRW=Yes|No

You can mount a read/write copy of your IBM Spectrum Protect backup so that you can modify the copy without invalidating the backup. You use this option to indicate whether a snapshot backup is mounted as read/write. The default value is as specified in the configuration file with the **/MOUNTRW** parameter. If a default value is not specified in the configuration file, the default value is No. By specifying the **/MOUNTRW** option, you override the default value.

The following values are available:

No

Perform a read-only mount operation.

Yes

Perform a read/write mount operation. The behavior of the read/write mount is controlled by the **USESNAPOFASNAPMount** parameter in the configuration file.

- If **USESNAPOFASNAPMount** is set to No, you can mount only COPY backups as read/write. After mounting, the original backup is modified and can no longer be used as a restore point in future database restore operations (on the **VSS Options** properties page, the **Mount read/write (modifies backup, applies to COPY backups only)** check box is selected).
- If **USESNAPOFASNAPMount** is set to Yes, you can mount both FULL and COPY backup types as read/write (on the **VSS Options** properties page, the **Mount read/write (without modifying backup)** check box is selected). In this instance, the backups are not modified and can be used in future restore operations.

Important:

This mount option is only available for the following devices:

- SAN Volume Controller (SVC) devices, which require IBM System Storage Support for Microsoft Virtual Disk and Volume Shadow Copy Services version 4.12 or later. Dynamic target allocation is not supported.
- XIV system devices, which require IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service version 2.9 or later.

You must allocate more target volumes on your storage device to accommodate the number of concurrent read/write mounts you want to do. An extra target volume that matches the size of the volume to be mounted, is needed for each concurrent read/write mount of that volume.

/PRESNapshotcmd=cmdstring

The **/presnapshotcmd** parameter runs a command or script before a snapshot operation begins. You can use this optional parameter to quiesce an application before a snapshot is created. You can then restart the application after the snapshot is started by using the **/postsnapshotcmd** optional parameter. The *cmdstring* variable specifies the command to run before the snapshot operation begins. You must specify the fully qualified path name for the command script.

/POSTSNapshotcmd=cmdstring

The **/postsnapshotcmd** parameter runs a command or script after a snapshot operation ends. You can use this optional parameter to resume the application after the snapshot is created. This parameter is used with the **/presnapshotcmd** parameter. The *cmdstring* variable must be a fully qualified path.

Update config example

This output example provides a sample of the text, messages, and process status that displays when you use the **update config** command.

The **fcmcli update config localdsmagentnode=server12** command sets the node name `server12` as the local system that performs the VSS backups. An example of the output is provided:

```
FMX5054I The preference has been set successfully.
```

The **fcmcli update config numberformat=2** command specifies that the 2 format is used to display numbers (n, nnn, dd.). An example of the output is provided:

```
FMX5054I The preference has been set successfully.
```

The **fcmcli update config localdsmagentnode=server44 / configfile=fcmcfg_server44.xml** command sets the node name `server44` as the local system that performs the VSS backups. This command also specifies that IBM Spectrum Protect Snapshot operations use the settings in the `fcmcfg_server44.xml` configuration file. An example of the output is provided:

```
FMX5054I The preference has been set successfully.
```

VSS policy commands

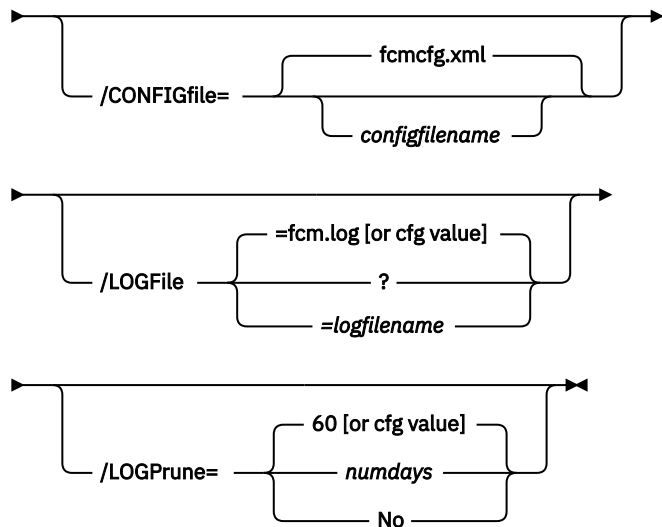
Use VSS policy commands to manage VSS policy binding statements.

UPDATE VSSPolicy

This command is used to update an existing VSS policy binding statement.

FCMCLI command: UPDATE VSSpolicy

```
➤ FCMCLI — UPDATE VSSpolicy — "policy_definition" — /SEQnumber= — number ➤
```



Parameters:

- **policy_definition**: Specifies the name of the VSS policy binding statement that is being updated.
- **SEQnumber**: Specifies the sequence priority for the updated policy binding statement.
- **CONFIGfile**: Specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use with the **update vsspolicy** command.

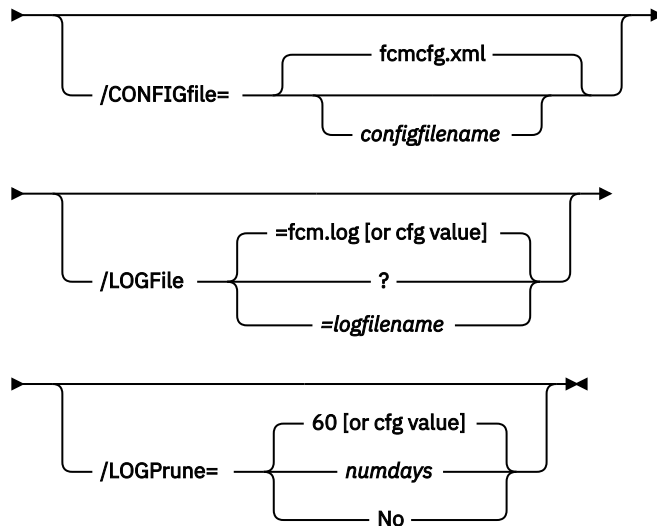
- **LOGFile:** Specify the name (*logfile*) of the activity log file to use with the **update vsspolicy** command.
- **LOGPrune:** Specify whether to disable log pruning or to prune the log for one command run. By default, log pruning occurs daily. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the pruning process.

INSert VSSpolicy

This command inserts a new VSS policy binding statement at the position that is specified by the **/SEQnumber** parameter.

FCMCLI command: INSert VSSpolicy

► FCMCLI — INSert VSSpolicy — "*policy_definition*" — /SEQnumber= — *number* —►



Parameters:

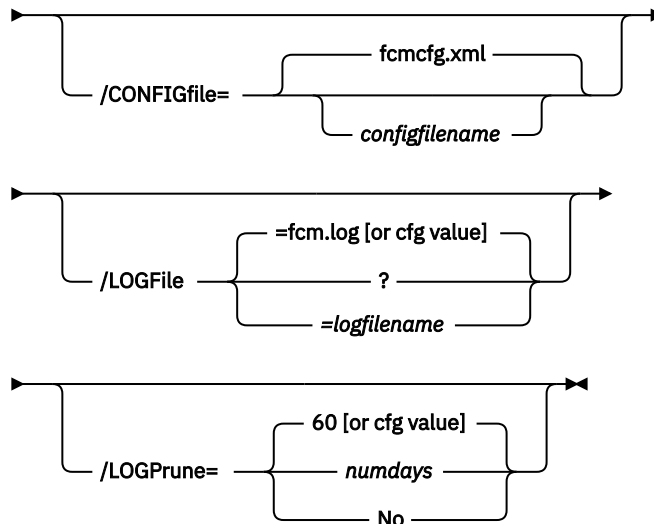
- *policy_definition*: Specifies the name of the VSS policy binding statement that is being updated.
- **SEQnumber**: Specifies the sequence priority of the inserted policy binding statement.
The default value is the sequence value of the highest prioritized VSS policy in the IBM Spectrum Protect Snapshot configuration file (*fcmcfg.xml*).
- **CONFIGfile**: Specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use with the **insert vsspolicy** command.
- **LOGFile**: Specify the name (*logfile*) of the activity log file to use with the **insert vsspolicy** command.
- **LOGPrune**: Specify whether to disable log pruning or to prune the log for one command run. By default, log pruning occurs daily. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the pruning process.

DELeTe VSSpolicy

This command is used to delete a VSS policy binding statement at the position that is specified by the **/SEQnumber** parameter.

FCMCLI command: DElete VSSpolicy

► FCMCLI — DElete VSSpolicy — /SEQnumber= — *number* →



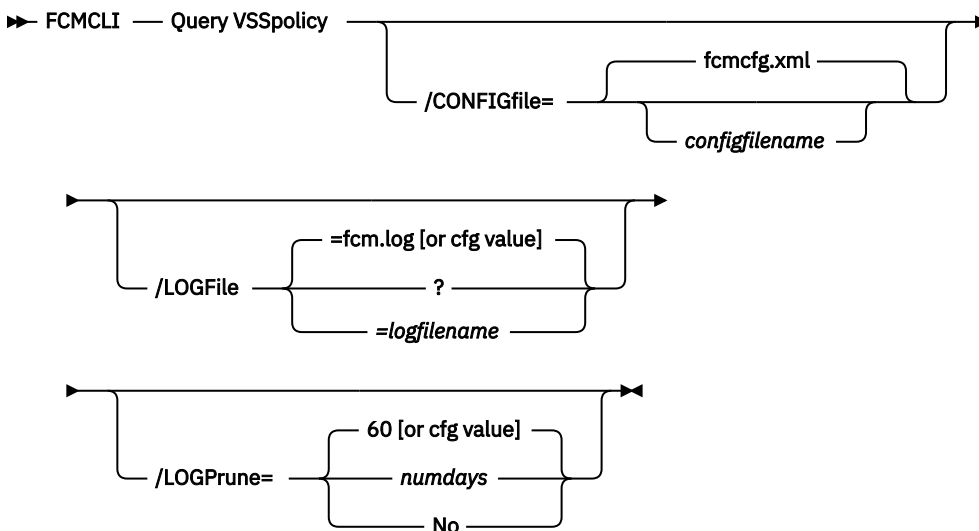
Parameters:

- **SEQnumber:** Specifies the sequence priority for the policy binding statement to delete.
- **CONFIGfile:** Specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file that contains the values to use with the **delete vsspolicy** command.
- **LOGFile:** Specify the name (*logfile*) of the activity log file to use with the **delete vsspolicy** command.
- **LOGPrune** Specify whether to disable log pruning or to prune the log for one command run. By default, log pruning occurs daily. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the pruning process.

Query VSSpolicy

This command is used to show the VSS policy binding statements in the configuration file.

FCMCLI command: Query VSSpolicy



Parameters:

- **CONFIGfile:** Specify the name (*configfilename*) of the IBM Spectrum Protect Snapshot configuration file to show.

- **LOGFile:** Specify the name (*logfile*) of the activity log file to use with the **query vsspolicy** command.
- **LOGPrune:** Specify whether to disable log pruning or to prune the log for one command run. By default, log pruning occurs daily. The *numdays* variable represents the number of days to save log entries. By default, 60 days of log entries are saved in the pruning process.

VSS policy command examples

The following output examples provide a sample of the text, messages, and process status that displays when you use the VSS policy commands.

In this example, the **fcmdi update vsspolicy "* * FULL LOCAL STANDARD" /SEQnumber=2** command updates the default VSS policy binding statement at sequence priority 2. The following output is displayed:

```
UPDATE VSSpolicy was successful.
```

In this example, the **fcmdi insert vsspolicy "* * FULL LOCAL STANDARD" /SEQnumber=2** command inserts the default VSS policy binding statement at sequence priority 2. The following output is displayed:

```
INSERT VSSpolicy was successful.
```

In this example, the **fcmdi delete vsspolicy /SEQnumber=1** command deletes the VSS policy binding statement at sequence priority 1. The following output is displayed:

```
DELETE VSSpolicy was successful.
```

In this example, the **fcmdi query vsspolicy /configfile=fcmcfg_server44.xml** command queries the VSS policy binding statements in the *fcmcfg_server44.xml* configuration file. The following output is displayed:

```
FCM for Windows VSS Policy
```

```
VSS policy statements are processed from the bottom up and processing stops at the first match. To ensure that more specific specifications are processed at all, the more general specification should be listed before the more specific ones, so as to be processed after the more specific specifications. Otherwise, the more general specification will match the target before the more specific specifications are seen.
```

```
-----
Sequence Number ..... 1
Server ..... SERVER44
Component ..... C:
Backup Type ..... FULL
Backup Destination ..... LOCAL
Management Class ..... STANDARD
```

Appendix A. Accessibility features for the IBM Spectrum Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Spectrum Protect family of products uses the latest W3C Standard, [WAI-ARIA 1.0](http://www.w3.org/TR/wai-aria/) (www.w3.org/TR/wai-aria/), to ensure compliance with [US Section 508](http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and [Web Content Accessibility Guidelines \(WCAG\) 2.0](http://www.w3.org/TR/WCAG20/) (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the [Accessibility](http://www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility) section of the [IBM Knowledge Center help](http://www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility) (www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Spectrum Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility](http://www.ibm.com/able) (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might 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, MD-NC119
Armonk, NY 10504-1785
US*

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*

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

Licensees 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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and 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 www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

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

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.
See the [IBM Spectrum Protect glossary](#).

Index

A

accessibility features [369](#)
active [282](#)
active parameter
 and query fcm command [215](#)
active/inactive state
 in restore operations [298](#)
all parameter
 and query backup command [342](#)
 and query fcm command [215](#)
AlwaysOn failover
 configure [56](#)
AlwaysOn node
 properties [56](#)
 transitioning standard databases to [75](#)
AlwaysOn preferences [56](#)
AlwaysOn scheduled backup preferences [56](#)
alwaysonpriority [262](#)
APAR [188](#)
authorization mode, setting
 using the CLI [265](#), [267](#)
auto select [132](#)
auto select option, GUI [102](#)
automated failover
 overview [25](#)
automated processing
 scripts [144](#)
automating
 custom applications and file system tasks tasks [168](#)
 Data Protection for Microsoft Exchange Server tasks [166](#)
 Data Protection for SQL Server tasks [167](#)
 tasks [166–168](#)
availability database restores
 overview [25](#), [121](#)

B

backing up custom application and file system data
[145–147](#)
backing up data [145–147](#)
backing up Exchange Server data [99](#)
backing up Exchange Server data in a DAG environment [99](#)
backing up SQL availability databases
 by using the legacy method [125](#)
 by using the VSS method [123](#)
backing up SQL databases
 by using the legacy method [125](#)
backing up SQL Server data
 by using the VSS method [123](#)
backup
 command line [198](#), [317](#), [350](#)
 copy-only full [119](#)
 file backup [119](#)
 full [93](#), [119](#)
 full plus differential plug log [119](#)
 full plus differentials [93](#), [119](#)

backup (*continued*)
 full plus incremental [93](#)
 full plus log [119](#)
 Legacy [9](#)
backup command
 and /BACKUPDESTination parameter [262](#)
 and /backupmethod parameter [263](#)
 and /FCMOPTFile parameter [342](#)
 and /logfile parameter [342](#), [352](#)
 and /OFFLOAD parameter [265](#)
 and /Quiet parameter [265](#)
 and /tsmnode parameter [343](#)
 and /tsmoptfile parameter [353](#)
 and /tsmpassword parameter [343](#), [354](#)
 overview [191](#), [316](#)
 syntax diagram [192](#), [317](#)
backup methods [7](#)
backup object types
 copyfull [262](#)
 full [262](#), [296](#)
backup operations
 time-saving strategy [119](#)
 using the GUI
 backup databases tab [123](#)
backup processing [4](#)
backup strategy
 copy-only full [119](#)
 file backup [119](#)
 full backup [93](#), [119](#)
 full plus differential plug log [119](#)
 full plus differentials [93](#), [119](#)
 full plus incremental [93](#)
 full plus log [119](#)
 group backup [119](#)
 IBM Spectrum Protect versus local shadow volumes [93](#),
 [119](#)
backupdestination [303](#)
BACKUPDESTination [193](#), [199](#), [250](#)
BACKUPDESTination parameter
 and backup command [262](#)
backupmethod parameter
 and backup command [263](#)
backups of availability databases
 overview [24](#)
binary sort order [287](#)

C

capacity
 determining managed storage [90](#)
capacity planning [27](#)
checksum
 SQL Server [123](#)
CHECKSum
 SQL Server [123](#)
CLIENTACcessserver [235](#)
clone operations [128](#)

- cluster [23](#)
- clustering
 - strategy [23](#)
- clustering state
 - querying [285](#)
- command line parameters
 - /all
 - and query backup [342](#)
 - and query fcm [215](#)
 - /backupmethod
 - and backup [263](#)
 - /configfile
 - and query fcm [215](#)
 - /OFFLOAD
 - and backup [265](#)
 - /pitdate
 - and mount backup [333](#), [358](#)
 - and restore [353](#)
 - /pittime
 - and mount backup [333](#), [358](#)
 - and restore [353](#)
 - /quiet
 - and restore [225](#)
 - /Quiet
 - and backup [265](#)
- command-line interface, IBM Spectrum Protect Snapshot
 - overview [259](#)
- command-line interface, IBM Spectrum Protect Snapshot
 - for custom applications
 - overview [316](#)
- command-line interface, IBM Spectrum Protect Snapshot
 - for Exchange Server
 - overview [190](#)
- command-line parameters
 - /active
 - and query fcm [215](#)
 - /BACKUPDESTINATION
 - and backup [262](#)
 - /configfile
 - and mount backup [331](#)
 - and query backup [342](#)
 - and query component [336](#)
 - and restore [222](#), [351](#)
 - and unmount backup [357](#)
 - /CONFIGfile
 - and query config [338](#)
 - /FCMOPTFile
 - and backup [342](#)
 - and mount backup [331](#)
 - and query component [336](#)
 - and query config [339](#)
 - /fromserver
 - and mount backup [332](#)
 - and query backup [342](#)
 - and restore [351](#)
 - /instantrestore
 - and restore [223](#), [351](#)
 - /into
 - and restore [351](#)
 - /intodb
 - and restore [223](#)
 - /logfile
 - and backup [342](#), [352](#)
 - and mount backup [275](#), [332](#)

- command-line parameters (*continued*)
 - /logfile (*continued*)
 - and query config [339](#)
 - and restore [223](#)
 - and unmount backup [357](#)
 - /LOGFile
 - and query component [336](#)
 - /mountdatabases
 - and restore [224](#)
 - /noprompt
 - and restore [352](#)
 - /object
 - and restore [225](#)
 - /recover
 - and restore [225](#)
 - /SHOWMAILBOXInfo
 - and query fcm [217](#)
 - /templogrestorepath
 - and restore [225](#)
 - /tsmnode
 - and backup [343](#)
 - and mount backup [207](#), [334](#)
 - and query component [337](#)
 - and query config [340](#)
 - and restore [353](#)
 - and unmount backup [258](#), [359](#)
 - /tsmoptfile
 - and mount backup [334](#)
 - and restore [353](#)
 - and unmount backup [359](#)
 - /tsmpassword
 - and backup [343](#)
 - and mount backup [334](#)
 - and query component [337](#)
 - and query config [340](#)
 - and restore [354](#)
 - and unmount backup [359](#)
 - dagnode [222](#)
- commands
 - capacity
 - management [217](#), [284](#), [346](#)
 - usage reports [217](#), [284](#), [346](#)
 - query config [338](#)
 - query fcm [213](#), [280](#)
 - query managedcapacity [217](#), [284](#), [346](#)
 - set [248](#)
 - update config [360](#)
- commands,
 - tdpsqlc set [307](#)
- commands, IBM Spectrum Protect Snapshot
 - query backup [341](#)
 - query component [335](#)
- commands, IBM Spectrum Protect Snapshot for Exchange Server
 - query exchange [210](#)
 - query tdp [218](#)
- commands, IBM Spectrum Protect Snapshot for SQL Server
 - query sql [285](#)
 - query tdp [291](#)
 - tdpsqlc help [272](#)
- communication protocol option [44](#)
- compatibility level
 - querying [285](#)
- compatibilityinfo [287](#)

- compression option [44](#)
- configfile [193](#), [205](#), [211](#), [218](#), [228](#), [254](#), [256](#), [270](#), [275](#), [287](#), [291](#), [296](#), [303](#), [311](#), [314](#), [322](#), [363](#)
- CONFIGfile [199](#), [235](#), [317](#)
- configfile parameter
 - and mount backup command [331](#)
 - and query backup command [342](#)
 - and query component command [336](#)
 - and query config command [338](#)
 - and query fcm command [215](#)
 - and restore command [222](#), [351](#)
 - and unmount backup command [357](#)
- CONFIGfile parameter [263](#)
- configuration
 - Mailbox Restore** [66](#)
 - options [43](#)
 - procedure
 - offloaded backups [69](#)
 - stand-alone snapshot support [58](#), [60](#)
 - wizard [58](#), [60](#), [61](#), [65](#), [66](#)
- configuration file [282](#)
- configuration file, IBM Spectrum Protect Snapshot for SQL
 - setting values, CLI [263](#)
- configuration files
 - non-default locations [46](#)
- configuration preferences [46](#)
- configuration settings [46](#)
- configuration tasks
 - IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core [76](#)
- configuration values
 - Windows Server Core [143](#)
- configure
 - AlwaysOn failover preferences [56](#)
 - AlwaysOn scheduled backup priority [56](#)
- configure AlwaysOn node [56](#)
- configure Data Protection for Exchange Server with DAG node [61](#), [65](#)
- configuring
 - backup priority of SQL replicas [75](#)
 - binding
 - policy [12](#)
 - IBM Spectrum Protect Snapshot [43](#)
 - IBM Spectrum Protect Snapshot with IBM Spectrum Protect [61](#), [65](#)
 - policy [12](#)
 - SQL Server clustered environments [70](#)
 - SQL Server environments [70](#)
 - where scheduled backups are run on SQL replicas [75](#)
- configuring IBM Spectrum Protect Snapshot for SQL Server for Windows Server configuration [80](#)
- configuring the backup-archive client
 - for Windows Server configuration on IBM Spectrum Protect Snapshot [79](#)
- consistency checker [123](#)
- continuous replication [8](#), [95](#)
- custom application and file system backups
 - deleting [148](#)
- custom application and file system data
 - overview [144](#)
 - planning [144](#)
 - restore considerations [4](#)
 - restoring
 - custom application and file system data [148](#)

- custom applications and file system tasks
 - automating [168](#)
- custom settings [55](#)

D

- DAG [15](#), [91](#), [104](#)
- DAG node
 - configuration [61](#), [65](#)
- dagnode [193](#), [199](#), [228](#), [236](#)
- dagnode parameter [222](#)
- data protection
 - Exchange with VSS backup-restore support
 - gathering information before calling IBM [186](#)
 - Exchange with VSS backup/restore support
 - general help [171](#)
 - general help [174](#)
 - troubleshooting [181](#)
- Data Protection for Microsoft Exchange Server
 - overview [91](#)
 - restore types [17](#)
- Data Protection for Microsoft Exchange Server tasks
 - automating [166](#)
- Data Protection for SQL
 - options file
 - clusternode [23](#)
- Data Protection for SQL Server
 - overview [119](#)
 - restore types [17](#)
 - security [119](#)
 - VSS instant restore [122](#)
- Data Protection for SQL Server tasks
 - automating [167](#)
- database
 - delete backup
 - command line [269](#)
 - restoring master [139](#)
 - restoring to alternate [137](#)
 - restoring with full-text catalogs and indexes [141](#)
- Database Availability Group
 - backup and restore [8](#), [95](#)
 - deployment example [84](#)
- database availability groups [91](#)
- database name
 - restorefiles
 - command line [303](#)
- database owner [133](#)
- dateformat [361](#)
- DATEformat [250](#)
- dateformat parameter [307](#)
- dbcc check options [123](#)
- deactivate operations
 - using the GUI [131](#)
- delete backup
 - database
 - command line [269](#)
- delete backup command
 - overview [196](#), [268](#), [320](#)
 - syntax diagram [198](#), [269](#), [321](#)
- deleting custom application and file system backups [148](#)
- deleting Exchange Server VSS backups [101](#)
- deleting SQL Server VSS backups [130](#)
- detail parameter [215](#)
- developerWorks wiki [188](#)

- diagnosing VSS issues [173](#)
- diagnostics properties [49](#), [185](#)
- differential restore
 - GUI [135](#)
- differential versus log backup strategy [119](#)
- disability [369](#)
- dsm.opt file
 - clusternode [44](#)
 - communication protocol [44](#)
 - compression [44](#)
 - enableclientencryptkey [45](#)
 - enablelanfree [45](#)
 - encryptiontype [45](#)
 - include.encrypt [45](#)
 - nodename [44](#)

E

- email support files [188](#)
- enableclientencryptkey option [45](#)
- enablelanfree option [45](#)
- ENABLEREplacementchars parameter [263](#)
- encryption [45](#)
- encryptiontype option [45](#)
- error log files [171](#), [172](#)
- example
 - query exchange command [212](#)
 - query fcm command [217](#)
 - query tdp command [220](#)
 - restorefiles command [306](#)
 - restoremailbox command [247](#)
 - set command [255](#)
 - update config command [365](#)
- Exchange backup
 - DAG environment [99](#)
 - VSS
 - GUI [99](#)
- Exchange Database Availability Group
 - managing with single policy [15](#)
- Exchange Server 2013, 2016, and 2019
 - requirements [92](#)
- Exchange Server VSS backup
 - deleting [101](#)
 - mounting [100](#)
- EXCLUDEDAGActive [193](#)
- EXCLUDEDAGPassive [194](#)
- EXCLUDEDDB [194](#)
- EXCLUDEdb parameter [263](#)
- EXCLUDEDUMPster [236](#)
- excludedumpster parameter [222](#)
- EXCLUDENONDAGDbs [194](#)
- expiring VSS backup s
 - policy [10](#)

F

- failover
 - overview [25](#)
- Failover clustering and AlwaysOn Availability [23](#)
- fcm.log file
 - and mount backup command [332](#)
 - and query backup command [342](#)
 - and query component command [336](#)

- fcm.log file (*continued*)
 - and query config command [339](#)
 - and restore command [352](#)
 - and unmount backup command [357](#)
- fcmcfg.xml file
 - and mount backup command [331](#)
 - and query backup command [342](#)
 - and query component command [336](#)
 - and query config command [338](#)
 - and restore command [351](#)
 - and unmount backup command [357](#)
- fcmcli.exe
 - overview [316](#)
- fcmoptfile [211](#), [236](#), [282](#), [296](#)
- FCMOPTFile [318](#)
- fcmoptfile parameter [215](#), [222](#)
- FCMOPTFile parameter
 - and backup command [342](#)
 - and mount backup command [331](#)
 - and query component command [336](#)
 - and query config command [339](#)
- file backup
 - strategy [119](#)
- files
 - dsm.opt [43](#)
 - fcm.log
 - and mount backup command [332](#)
 - and query backup command [342](#)
 - and query component command [336](#)
 - and query config command [339](#)
 - and restore command [352](#)
 - and unmount backup command [357](#)
 - fcmcfg.xml
 - and mount backup command [331](#)
 - and query backup command [342](#)
 - and query component command [336](#)
 - and query config command [338](#)
 - and restore command [351](#)
 - and unmount backup command [357](#)
 - fcmcli.exe [316](#)
 - IBM Spectrum Protect Snapshot options [334](#), [337](#), [340](#), [343](#), [354](#), [359](#)
 - tdpexc.cfg
 - and query fcm command [215](#)
 - and restore command [222](#), [225](#)
 - tdpexc.log
 - and query fcm command [216](#)
 - and restore command [223](#)
 - tdpexcc.exe [190](#)
 - tdpsql.log
 - and mount backup command [275](#)
 - tdpsqlc.exe [259](#)
- FlashCopy Manager for Microsoft Exchange ServerIBM Spectrum Protect Snapshot for Microsoft SQL Server
 - policy settings [10](#)
- from server option, GUI [102](#)
- FROMArchive [236](#)
- fromexcserver [199](#), [205](#), [215](#), [223](#), [228](#)
- fromserver [322](#)
- fromserver parameter
 - and mount backup command [332](#)
 - and query backup command [342](#)
 - and restore command [351](#)

- fromsqlserver [270](#), [275](#), [283](#), [296](#), [303](#), [308](#)
- full backup
 - strategy [93](#), [119](#)
- full plus differential backup
 - strategy [93](#), [119](#)
- full plus differential plus log backup
 - strategy [119](#)
- full plus incremental backup
 - strategy [93](#)
- full plus log backup
 - strategy [119](#)
- full restore
 - GUI [135](#)

G

- general properties for Exchange Server [51](#)
- general properties for SQL Server [50](#)
- graphical user interface (GUI)
 - backup databases tab [123](#)
 - inactivating SQL databases [131](#)
 - restore options [102](#)
- group backup
 - strategy [119](#)
- group restore
 - GUI [135](#)
- GUI
 - DAG Exchange backup [99](#)
 - Exchange VSS backup [99](#)
 - individual mailbox restore [107](#)
 - restore options [103](#)
- GUI)
 - restore options [132](#)
- guidelines
 - VSS restore [4](#)

H

- help command
 - syntax diagram [202](#), [272](#), [324](#)
- help command, IBM Spectrum Protect Snapshot for SQL Server
 - described [272](#)

I

- IBM Knowledge Center [xi](#)
- IBM SAN Volume Controller and IBM Storwize V7000
 - using IBM Spectrum Protect Snapshot [21](#), [86](#), [174](#)
- IBM Spectrum Protect
 - configuring options [43](#)
 - policy settings [10](#)
- IBM Spectrum Protect **Configuration**
 - wizard [61](#)
- IBM Spectrum Protect server scheduler [23](#)
- IBM Spectrum Protect Snapshot
 - backups [7](#)
 - commands [259](#)
 - commands for custom applications [316](#)
 - configuring [43](#)
 - install [30](#)
 - install prerequisites [29](#)
 - install, upgrade, migrate [29](#)

- IBM Spectrum Protect Snapshot (*continued*)
 - migrate [38](#)
 - migrate DAG backups [38](#)
 - overview [1](#)
 - planning [27](#)
 - reference [189](#)
 - silent installation on Windows Server Core (MSI) [37](#)
 - silent installation on Windows Server Core (spinstall.exe) [36](#)
 - upgrade [37](#)
 - VSS configuring [67](#)
 - VSS planning [2](#)
- IBM Spectrum Protect Snapshot for Exchange Server
 - commands [190](#)
- IBM Spectrum Protect Snapshot for SQL
 - configuration file, setting
 - using the CLI [263](#)
- IBM Spectrum Protect Snapshot for SQL Server
 - running on Windows Failover Cluster [23](#)
- IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core
 - overview [141](#)
- IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core
 - configuration tasks [76](#)
- IBM Spectrum Protect Snapshot for Windows
 - creating an installation package on a DVD [35](#)
 - silent installation [32](#)
 - silent installation with batch file [36](#)
- IBM Spectrum Protect Snapshot for Windows silent installation
 - playing back the installation [35](#)
 - setup error messages [36](#)
- IBM Spectrum Protect Snapshot GUI
 - protecting [89](#)
 - starting [89](#)
- IBM Spectrum Protect Snapshot on Windows Server Core
 - silent installation [36](#)
- IBM Spectrum Protect Snapshot scripts
 - adding [187](#)
 - editing [187](#)
 - viewing [187](#)
- IBM Spectrum Protect SnapshotVSS backup
 - policy binding [13](#)
- IBM System Storage DS8000 series
 - requirements [20](#)
- IMPORTVSSNAPSHOTSONLY Whenneeded [251](#), [308](#), [361](#)
- include.encrypt option [45](#)
- IncludeTsmVm [134](#)
- indexes and tables
 - backing up [119](#)
- individual mailbox
 - restoremailbox
 - command line [235](#)
- individual mailbox restore
 - GUI [107](#)
- install
 - IBM Spectrum Protect Snapshot [30](#)
- install prerequisites
 - IBM Spectrum Protect Snapshot [29](#)
- install, upgrade, migrate
 - IBM Spectrum Protect Snapshot [29](#)
- installation
 - configuring options [43](#)

- installing
 - creating an installation package on a DVD [35](#)
 - silently with batch file [36](#)
 - silently with msixec.exe [37](#)
- installing IBM Spectrum Protect Snapshot
 - on multiple servers (silent) [32](#)
 - unattended (silent) [32](#)
- instant restore [135](#)
- instantrestore [296](#)
- instantrestore parameter
 - and restore command [223](#), [351](#)
- integrated user ID mode [265](#)
- into [228](#), [297](#), [305](#)
- into parameter
 - and restore command [351](#)
- intodb parameter
 - and restore command [223](#)

K

- keep cdc [133](#)
- keepcdc [297](#)
- KEEPRDB [236](#)
- keyboard [369](#)
- Knowledge Center [xi](#)

L

- language [362](#)
- Legacy backup
 - overview [9](#)
- local backup policy
 - setting [12](#)
- localdsmagentnode [309](#), [362](#)
- LOCALDSMAgentnode [251](#)
- log backup
 - strategy [123](#)
- log files
 - using for problem determination [171](#), [172](#)
- log restore
 - GUI [135](#)
- logfile [200](#), [205](#), [211](#), [219](#), [228](#), [257](#), [270](#), [283](#), [287](#), [297](#), [304](#), [309](#), [314](#), [362](#)
- LOGFile [194](#), [236](#), [251](#), [318](#), [322](#)
- logfile parameter
 - and mount backup command [275](#), [332](#)
 - and query backup command [342](#)
 - and query config command [339](#)
 - and query fcm command [216](#)
 - and restore command [223](#), [352](#)
 - and unmount backup command [357](#)
- LOGFile parameter
 - and query component command [336](#)
- logging properties [52](#)
- login settings
 - using the CLI [265](#)
- logprune [194](#), [200](#), [206](#), [219](#), [229](#), [237](#), [251](#), [270](#), [275](#), [283](#), [287](#), [292](#), [297](#), [304](#), [309](#), [314](#), [318](#), [323](#), [362](#)
- logprune parameter [212](#), [216](#), [224](#), [257](#), [264](#), [332](#), [337](#), [339](#), [342](#), [352](#), [357](#)

M

- mailbox
 - restoremailbox
 - command line [235](#)
 - mailbox history handling [180](#)
 - mailbox restore
 - guidelines [93](#)
 - overview [18](#)
 - requirements [92](#)
- Mailbox Restore Only Configuration**
 - wizard [66](#)
- mailbox restore operations
 - permissions [91](#)
 - security [91](#)
- mailboxfilter [237](#)
- mailboxoriglocation [240](#)
- mailboxrestoredate [240](#)
- mailboxrestoredestination [241](#)
- mailboxrestoretime [241](#)
- MAILBOXRestoreunread [245](#)
- managed storage
 - determining capacity [90](#)
- managing with single policy
 - Exchange Database Availability Group [15](#)
- MAPI
 - ensuring successful connections [97](#)
- MAPI settings for Exchange Server [57](#), [97](#)
- master database, restoring [139](#)
- migrate
 - DAG backups [38](#)
 - IBM Spectrum Protect Snapshot [38](#)
- migration
 - mailbox history handling [180](#)
- MINimumbackupinterval [195](#)
- MMC GUI
 - starting [89](#)
- mount backup command
 - and /configfile parameter [331](#)
 - and /FCMOPTFile parameter [331](#)
 - and /fromserver parameter [332](#)
 - and /logfile parameter [275](#), [332](#)
 - and /pitdate parameter [333](#), [358](#)
 - and /pittime parameter [333](#), [358](#)
 - and /tsmnode parameter [207](#), [334](#)
 - and /tsmoptfile parameter [334](#)
 - and /tsmpassword parameter [334](#)
 - overview [203](#), [273](#), [329](#)
 - syntax diagram [203](#), [273](#), [329](#)
- mountdatabases parameter
 - and restore command [224](#)
- mounting Exchange Server VSS backups [100](#)
- mountrw [206](#), [245](#), [252](#), [254](#), [311](#), [332](#), [364](#)
- mountwait [195](#), [224](#), [229](#), [304](#)
- MOUNTWait [252](#)
- msixec.exe
 - used for silent installation [37](#)
- multiple-user mode [132](#)

N

- New in IBM Spectrum Protect Snapshot Version 8.1.7 on Windows [xv](#)
- node name

node name (*continued*)
 offloaded backup [68](#)
 proxy nodes [67](#)
 VSS [67](#)
nodename option [44](#)
noprompt parameter
 and restore command [352](#)
numberformat [310](#), [363](#)
NUMberformat [252](#)

O

object [201](#), [230](#), [271](#), [284](#), [298](#), [305](#)
object parameter
 and restore command [225](#)
OFFLOAD [195](#), [318](#)
OFFLOAD parameter
 and backup command [265](#)
offloaded backup
 configuration procedure [69](#)
 node names [68](#)
offloaded VSS backup
 overview [8](#)
olderthan [201](#)
optional parameters [228](#), [296](#), [303](#)
options
 GUI restore
 mountdatabases [103](#)
 run recovery [103](#)
options file, Data Protection for SQL
 cluster [23](#)
options files
 non-default locations [46](#)
overview
 availability database restores [25](#), [121](#)
 backups of availability databases [24](#)
 IBM Spectrum Protect Snapshot for SQL Server on
 Windows Serer Core [141](#)
 Legacy backup [9](#)
 offloaded VSS backup [8](#)
 thin provisioning support [22](#)
 VSS backup [8](#)

P

parameters
 /active
 and query fcm command [215](#)
 /all
 and query backup command [342](#)
 and query fcm command [215](#)
 /BACKUPDESTination
 and backup command [262](#)
 /backupmethod
 and backup command [263](#)
 /configfile
 and mount backup command [331](#)
 and query backup command [342](#)
 and query component command [336](#)
 and query fcm command [215](#)
 and restore command [222](#), [351](#)
 and unmount backup command [357](#)
 /CONFIGfile

parameters (*continued*)
 /CONFIGfile (*continued*)
 and query config command [338](#)
 /FCMOPtFile
 and backup command [342](#)
 and mount backup command [331](#)
 and query component command [336](#)
 and query config command [339](#)
 /fromserver
 and mount backup command [332](#)
 and query backup command [342](#)
 and restore command [351](#)
 /instantrestore
 and restore command [223](#), [351](#)
 /into
 and restore command [351](#)
 /intodb
 and restore command [223](#)
 /logfile
 and mount backup command [275](#), [332](#)
 and query backup command [342](#)
 and query config command [339](#)
 and query fcm command [216](#)
 and restore command [223](#), [352](#)
 and unmount backup command [357](#)
 /LOGFile
 and query component command [336](#)
 /mountdatabases
 and restore command [224](#)
 /noprompt
 and restore command [352](#)
 /object
 and restore command [225](#)
 /OFFLOAD
 and backup command [265](#)
 /pitdate
 and mount backup command [333](#), [358](#)
 and restore command [353](#)
 /pittime
 and mount backup command [333](#), [358](#)
 and restore command [353](#)
 /quiet
 and restore command [225](#)
 /Quiet
 and delete backup command [265](#)
 /recover
 and restore command [225](#)
 /SHOWMAILBOXInfo
 and query fcm command [217](#)
 /templogrestorepath
 and restore parameter [225](#)
 /tsmnode
 and backup command [343](#)
 and mount backup command [207](#), [334](#)
 and query component command [337](#)
 and query config command [340](#)
 and restore command [353](#)
 and unmount backup command [258](#), [359](#)
 /tsmoptfile
 and mount backup command [334](#)
 and restore command [353](#)
 and unmount backup command [359](#)
 /tsmpassword
 and backup command [343](#)

parameters (*continued*)

/tsmpassword (*continued*)

- and mount backup command [334](#)
- and query component command [337](#)
- and query config command [340](#)
- and restore command [354](#)
- and unmount backup command [359](#)

dagnode [222](#)

set command [307](#)

parameters, described

optional

- /CONFIGfile [263](#)
- /detail [215](#)
- /EXCLUDEdb [263](#)
- /EXCLUDEDUMPster [222](#)
- /fcmoptfile [222](#)
- /FCMOPTFile [215](#), [264](#)
- /LOGFile [264](#)
- /logprune [212](#), [216](#), [224](#), [257](#), [264](#), [332](#), [337](#), [339](#), [342](#), [352](#), [357](#)
- /SQLAUTHentication [265](#)
- /SQLPassword [265](#)
- /SQLUser [267](#)
- ENBLEReplacementchars [263](#)
- SQLCHECKSum [265](#), [288](#)

positional

copyfull [262](#)

performance properties [55](#)

pitdate [276](#)

pitdate parameter

- and mount backup command [333](#), [358](#)
- and restore command [353](#)

pittime [276](#)

pittime parameter

- and mount backup command [333](#), [358](#)
- and restore command [353](#)

planning [27](#)

policy

- binding [12](#)
- binding IBM Spectrum Protect Snapshot VSS backups [13](#)
- configuring [12](#)
- expiring VSS backup s [10](#)
- setting local policy [12](#)

policy management properties [48](#)

policy settings

- IBM Spectrum Protect Snapshot for Microsoft Exchange Server
- IBM Spectrum Protect Snapshot for Microsoft SQL Server and IBM Spectrum Protect [10](#)

postsnapshot command [52](#)

postsnapshotcmd [145–147](#), [364](#)

POSTSNapshotcmd [319](#)

pre/post snapshot properties [52](#)

preferdagspassive [195](#)

preferences [47](#)

presnapshot command [52](#)

presnapshotcmd [145–147](#), [364](#)

PRESNapshotcmd [319](#)

printing reports [158](#)

problem determination [171](#), [172](#)

product support [188](#)

properties

- AlwaysOn node [56](#)
- AlwaysOn preferences [56](#)

properties (*continued*)

custom settings [55](#)

diagnostics [49](#), [185](#)

general Exchange Server [51](#)

general SQL Server [50](#)

logging [52](#)

MAPI settings [57](#), [97](#)

performance [55](#)

policy management [48](#)

pre/post snapshot [52](#)

regional settings [53](#)

SQL login [50](#)

VSS [53](#)

property pages [47](#)

protecting

IBM Spectrum Protect Snapshot GUI [89](#)

protecting SQL Server data with IBM Spectrum Protect Snapshot

on Windows Server Core [141](#)

Windows Sever Core

- protecting SQL Server data with IBM Spectrum Protect Snapshot [141](#)

proxy nodes [67](#)

publications [xi](#)

Q

query backup command

- and /all parameter [342](#)
- and /configfile parameter [342](#)
- and /fromserver parameter [342](#)

query backup command, IBM Spectrum Protect Snapshot

overview [341](#)

syntax diagram [341](#)

query component command

- and /configfile parameter [336](#)
- and /FCMOPTFile parameter [336](#)
- and /LOGFile parameter [336](#)
- and /tsmnode parameter [337](#)
- and /tsmpassword parameter [337](#)
- syntax diagram [336](#)

query component command, IBM Spectrum Protect Snapshot

overview [335](#)

query config command

- and /CONFIGfile parameter [338](#)
- and /FCMOPTFile parameter [339](#)
- and /logfile parameter [339](#)
- and /tsmnode parameter [340](#)
- and /tsmpassword parameter [340](#)
- overview [338](#)

syntax diagram [338](#)

query exchange command

- example [212](#)
- syntax diagram [211](#), [286](#)

query exchange command, IBM Spectrum Protect Snapshot for Exchange Server

overview [210](#)

query fcm command

- and /active parameter [215](#)
- and /all parameter [215](#)
- and /configfile parameter [215](#)
- and /logfile parameter [216](#)
- and /SHOWMAILBOXInfo parameter [217](#)

- query fcm command (*continued*)
 - example [217](#)
 - overview [213](#), [280](#)
 - syntax diagram [214](#)
- query fcm command, IBM Spectrum Protect Snapshot for SQL Server
 - syntax diagram [281](#)
- query operations
 - query SQL [286](#)
- query sql command, IBM Spectrum Protect Snapshot for SQL Server
 - overview [285](#)
- query tdp command
 - example [220](#)
- query tdp command, IBM Spectrum Protect Snapshot for Exchange Server
 - overview [218](#)
 - syntax diagram [218](#), [291](#)
- query tdp command, IBM Spectrum Protect Snapshot for SQL Server
 - overview [291](#)
- quiet [201](#), [230](#), [246](#), [271](#), [298](#), [305](#)
- Quiet [195](#)
- quiet parameter
 - and restore command [225](#)
- Quiet parameter
 - and delete backup command [265](#)

R

- RBAC
 - permissions [91](#)
- reconciliation [144](#)
- recover parameter
 - and restore command [225](#)
- recovery [133](#), [298](#)
- reference
 - IBM Spectrum Protect Snapshot [189](#)
- regional properties [53](#)
- registering a client node
 - for Windows Server Core configuration on IBM Spectrum Protect Snapshot [77](#)
- relocatedir [299](#)
- remote diagnostics [185](#)
- remote system
 - stand-alone snapshot support [60](#)
- remote system configuration [65](#)
- replace [133](#)
- replay option, GUI [103](#)
- replication copies [8](#), [95](#)
- reports
 - viewing, printing, and saving [158](#)
- requirements
 - IBM System Storage DS8000 series [20](#)
 - SAN Volume Controller [20](#)
 - Storwize family [20](#)
 - XIV [20](#)
- restore
 - database [17](#)
 - mailbox [18](#)
 - restorefiles command [17](#)
 - stripes [133](#)
 - transaction log [17](#)
 - types [17](#)

- restore command
 - and /configfile parameter [222](#), [351](#)
 - and /fromserver parameter [351](#)
 - and /instantrestore parameter [223](#), [351](#)
 - and /into parameter [351](#)
 - and /intodb parameter [223](#)
 - and /logfile parameter [223](#)
 - and /mountdatabases parameter [224](#)
 - and /noprompt parameter [352](#)
 - and /object parameter [225](#)
 - and /pitdate parameter [353](#)
 - and /pittime parameter [353](#)
 - and /quiet parameter [225](#)
 - and /recover parameter [225](#)
 - and /templogrestorepath parameter [225](#)
 - and /tsmnode parameter [353](#)
 - overview [220](#), [349](#)
 - syntax diagram [221](#), [350](#)
- Restore command [296](#)
- restore considerations
 - custom application and file system data [4](#)
- restore guidelines
 - SQL VSS [4](#)
- restore operations
 - GUI [132](#)
 - master database [139](#)
 - of inactive objects [298](#)
 - SQL databases with full-text catalogs and indexes [141](#)
 - to alternate instance [137](#)
 - using the GUI
 - auto select option [102](#)
 - from server option [102](#)
 - instant restore [103](#)
 - replay option [103](#)
 - restore options [102](#)
- restore options
 - GUI
 - mountdatabases [103](#)
 - run recovery [103](#)
- restore SQL
 - Windows Server Core [142](#)
- restoredate [299](#)
- restorefiles
 - command [227](#)
 - parameters [227](#)
- restorefiles command
 - backups [226](#)
 - example [306](#)
 - overview [226](#)
 - syntax diagram [227](#)
 - VSS [226](#)
- restoremailbox
 - individual mailbox
 - command line [235](#)
 - mailbox
 - command line [235](#)
 - syntax diagram [233](#)
- restoremailbox command
 - example [247](#)
 - overview [230](#)
- restoretme [299](#)
- restoring data
 - Exchange Server 2010 [112](#)
 - Exchange Server 2013 [112](#)

- restoring data (*continued*)
 - mailbox [107](#)
 - Mailbox Restore Browser [112](#)
- restoring SQL availability databases [135](#)
- restoring VE databases [135](#)
- restoring custom application and file system data [150](#)
- Role Based Access Control
 - permissions [91](#)

S

- sample output
 - set command [312](#)
- saving reports [158](#)
- scheduled backups
 - AlwaysOn preferences [56](#)
- scheduling backups [23](#)
- scheduling tasks [169](#)
- scripts for automated processing [144](#)
- security [119](#)
- security requirements [91](#)
- sending support files by using email [188](#)
- server, SQL
 - querying [286](#)
- Service Management Console [188](#)
- set
 - positional parameters [307](#)
- set command
 - example [255](#)
 - overview [248](#)
 - sample output [312](#)
 - syntax diagram [249](#)
 - Windows Server Core [143](#)
- set restore
 - GUI [135](#)
- setting up a proxy node for offloaded VSS backups
 - for Windows Server configuration on IBM Spectrum Protect Snapshot [78](#)
- SHOWMAILBOXInfo parameter
 - and query fcm command [217](#)
- silent installation
 - IBM Spectrum Protect Snapshot on Windows Server Core [36](#)
 - with spinstall.exe [32](#)
- silent installation (MSI)
 - IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core [37](#)
- silent installation (spinstall.exe)
 - IBM Spectrum Protect Snapshot for SQL Server on Windows Server Core [36](#)
- silent installation of IBM Spectrum Protect Snapshot [32](#)
- single-user mode [132](#)
- SKIPINTEGRITYCHECK [195](#)
- space, saving
 - strategy [119](#)
- spinstall.exe
 - used for silent installation [32](#)
- SQL 2012 Server
 - running in cluster environment [23](#)
- SQL availability databases
 - backing up with the legacy method [125](#)
 - backing up with VSS [123](#)
 - restoring [135](#)
- SQL database cloning

- SQL database cloning (*continued*)
 - using Windows PowerShell cmdlet [128](#)
- SQL databases with full-text catalogs and indexes, restoring [141](#)
- SQL login properties [50](#)
- SQL server
 - user ID [265](#)
- SQL Server 2000
 - differential strategy [119](#)
- SQL Server 2012 services [23](#)
- SQL Server integrity checking
 - checksum [123](#)
- SQL Server VSS backup
 - deleting [130](#)
- SQL VSS restore considerations [4](#)
- sqlauthentication [288](#), [299](#), [310](#)
- SQLAUTHentication parameter [265](#)
- SQLCHECKSum [312](#)
- SQLCHECKSum parameter [265](#), [288](#)
- sqlcompression [312](#)
- sqlpassword [288](#), [299](#)
- SQLPassword parameter [265](#)
- sqlserver [299](#), [310](#)
- sqlserver parameter [265](#), [288](#)
- sqluser [290](#), [301](#)
- SQLUser parameter [267](#)
- sqluserid parameter [310](#)

Standalone Configuration

- wizard [58](#), [60](#)
- standby server undo file [134](#)
- starting
 - IBM Spectrum Protect Snapshot GUI [89](#)
 - MMC GUI [89](#)
- storage
 - determining managed capacity [90](#)
- storage group
 - restorefiles
 - command line [227](#)
 - VSS backup
 - GUI [99](#)
- storage management, policy [10](#)
- Storwize family
 - requirements [20](#)
- syntax diagrams
 - backup command [192](#), [317](#)
 - delete backup command [198](#), [269](#), [321](#)
 - help command [202](#), [272](#), [324](#)
 - mount backup command [203](#), [273](#), [329](#)
 - query backup command, IBM Spectrum Protect Snapshot [341](#)
 - query component command [336](#)
 - query config command [338](#)
 - query exchange command [211](#), [286](#)
 - query fcm command [214](#)
 - query tdp command, IBM Spectrum Protect Snapshot for Exchange Server [218](#)
 - query tdp command, IBM Spectrum Protect Snapshot for SQL Server [291](#)
 - restore command [221](#), [350](#)
 - restorefiles command [227](#)
 - restoremailbox [233](#)
 - set command [249](#)
 - unmount backup command [256](#), [313](#), [356](#)
 - update config command [360](#)

syntax diagrams, IBM Spectrum Protect Snapshot for SQL Server
query fcm command [281](#)
sysadmin fixed server role [293](#), [310](#)

T

tables and indexes
backing up [119](#)
tasks
automating [166](#)–[168](#)
tdpexc.cfg file
and query fcm command [215](#)
and restore command [222](#), [225](#)
tdpexc.log file
and query fcm command [216](#)
and restore command [223](#)
tdpexcc.exe
overview [190](#)
tdpsql.cfg, setting values
using the CLI [263](#)
tdpsql.log file
and mount backup command [275](#)
tdpsqlc.exe
overview [259](#)
tempdb [124](#)
tempdbrestorepath [246](#)
TEMPDBRESTorepath [253](#)
templogrestorepath [246](#)
templogrestorepath parameter
and restore command [225](#)
tempmailboxalias [247](#)
timeformat [310](#), [363](#)
TIMEformat [253](#)
Tivoli Storage FlashCopy Manager
backups [4](#)
Tivoli Storage FlashCopy Manager trace and log files
gathering [183](#), [185](#)
trace and log files
gathering [183](#)
transact-SQL command [132](#)
transaction log
restore [17](#), [220](#)
transitioning standard SQL databases to the AlwaysOn node
[75](#)
troubleshooting
VSS issues [173](#)
TSM Configuration
wizard [65](#)
tsmnode [201](#), [277](#), [305](#), [315](#)
TSMNODE [319](#), [323](#)
tsmnode parameter
and backup command [343](#)
and mount backup command [207](#), [334](#)
and query component command [337](#)
and query config command [340](#)
and restore command [353](#)
and unmount backup command [258](#), [359](#)
tsmoptfile [207](#), [230](#), [258](#), [277](#), [305](#), [315](#)
TSMOPTFile [196](#)
tsmoptfile parameter
and mount backup command [334](#)
and restore command [353](#)
and unmount backup command [359](#)

tsmpassword [207](#), [258](#), [277](#), [306](#), [315](#)
TSMPassword [319](#), [323](#)
tsmpassword parameter
and backup command [343](#)
and mount backup command [334](#)
and query component command [337](#)
and query config command [340](#)
and restore command [354](#)
and unmount backup command [359](#)

U

uninstalling IBM Spectrum Protect Snapshot [39](#)
unmount backup command
and /configfile parameter [357](#)
and /logfile parameter [357](#)
and /tsmnode parameter [258](#), [359](#)
and /tsmoptfile parameter [359](#)
and /tsmpassword parameter [359](#)
overview [256](#), [313](#), [356](#)
syntax diagram [256](#), [313](#), [356](#)
update config command
example [365](#)
overview [360](#)
syntax diagram [360](#)
UpdateMailboxInfoOnly [196](#)
upgrade
IBM Spectrum Protect Snapshot [37](#)
usealwaysonnode [267](#)
USEEXISTINGRDB [247](#)
user mode, setting [132](#)
USESNAPOFASNAPTomount [254](#)
using command-line help
Windows Server Core [90](#)
using IBM Spectrum Protect Snapshot
with IBM SAN Volume Controller and IBM Storwize
V7000 [21](#), [86](#), [174](#)

V

Verify Only option [134](#)
verifyonly [301](#)
viewing reports [158](#)
viewing system information for IBM Spectrum Protect
Snapshot for Windows [187](#)
VSS
node names [67](#)
overview [2](#)
proxy nodes [67](#)
restore into alternate locations [18](#)
VSS backup
characteristics [2](#)
overview [8](#)
planning requirements [3](#)
policy binding [13](#)
VSS configuring [67](#)
VSS fast restore
method [17](#)
VSS instant restore
method [17](#)
VSS planning [2](#)
VSS properties [53](#)
VSS provider [2](#)

- VSS Requestor [2](#)
- VSS restore
 - characteristics [3](#)
- VSS writer [2](#)
- VSSPOLICY, statements [14](#)

W

- wait for tape mounts [134](#)
- Wait for Tape Mounts for File Information [134](#)
- Windows authentication mode, setting
 - using the CLI [265](#)
- Windows Failover Cluster environment
 - converting to [23](#)
 - supported [23](#)
- Windows Server Core
 - backing up [141](#)
 - command-line help [90](#)
 - restore SQL [142](#)
- Windows Server Core configuration
 - configuring IBM Spectrum Protect Snapshot for SQL Server [80](#)
- Windows Server Core configuration on IBM Spectrum Protect Snapshot
 - configuring the backup-archive client [79](#)
 - registering a client node [77](#)
 - setting up a proxy node for offloaded backups [78](#)

X

- XIV
 - requirements [20](#)



Product Number: 5725-X22
5608-AB8