

IBM Spectrum Protect

*Blueprint and Server Automated
Configuration for Linux x86*

Version 5.1



Note:

Before you use this information and the product it supports, read the information in [“Notices”](#) on page [83](#).

First edition (September 2022)

This edition applies to Version 8.1.16 and later of the IBM Spectrum® Protect server, and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

© **Copyright International Business Machines Corporation 2013, 2022.**

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

Contents

About this document.....	v
Support for IBM Spectrum Protect blueprint and server automated configuration.....	v
What's new in Version 5.1.....	vii
Chapter 1. Introduction.....	1
Chapter 2. Implementation requirements.....	3
Hardware and software prerequisites.....	5
Hardware requirements.....	5
Software requirements.....	10
Planning worksheets.....	10
Chapter 3. Storage configuration blueprints.....	15
Small FlashSystem configuration.....	17
Medium FlashSystem configuration.....	18
Large FlashSystem configuration.....	20
IBM Elastic Storage System.....	21
Chapter 4. Setting up the system.....	23
Step 1: Set up and configure hardware.....	23
Step 2: Install the operating system.....	25
Step 3: IBM FlashSystem Storage: Configure multipath I/O.....	28
Step 4: IBM FlashSystem Storage: Configure file systems for IBM Spectrum Protect.....	29
Configure a file system by using the script.....	30
Configure a file system by using the manual procedure.....	30
Step 5: IBM Elastic Storage System: Configuring the system.....	31
Step 6: Test system performance.....	34
Step 7: Install the IBM Spectrum Protect backup-archive client.....	37
Step 8: Install the IBM Spectrum Protect server.....	38
Obtain the installation package.....	38
Install the IBM Spectrum Protect server.....	39
Chapter 5. Configuring the IBM Spectrum Protect server.....	41
Removing an IBM Spectrum Protect blueprint configuration.....	48
Chapter 6. Completing the system configuration.....	49
Changing initial passwords.....	49
Registering nodes and associating them with predefined client schedules.....	49
Reorganizing database tables and indexes.....	51
Chapter 7. Next steps.....	53
Optional: Set up data replication by using replication storage rules and subrules.....	53
Appendix A. Performance results.....	57
Extra Small system performance measurements.....	57
Small system performance measurements.....	58
Medium system performance measurements.....	59
Large system performance measurements.....	60

Workload simulation tool results.....	61
Appendix B. Configuring the disk system by using commands.....	65
Appendix C. Using a response file with the Blueprint configuration script.....	71
Appendix D. Using predefined client schedules.....	73
Appendix E. Modification of blueprint configurations.....	77
Appendix F. Troubleshooting.....	79
Appendix G. Accessibility.....	81
Notices.....	83
Index.....	87

About this document

This information is intended to facilitate the deployment of an IBM Spectrum Protect server by using detailed hardware specifications to build a system and automated scripts to configure the software. To complete the tasks, you must have an understanding of IBM Spectrum Protect and scripting.

Support for IBM Spectrum Protect blueprint and server automated configuration

The information in this document is distributed on an "as is" basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Spectrum Protect support is entitled and where the issues are not specific to a blueprint implementation.

What's new in Version 5.1

The IBM Spectrum Protect Blueprint configuration script, hardware and software requirements, and documentation are updated.

Replicate data to multiple servers

IBM Spectrum Protect 8.1.13 introduces multi-target replication, which is the process of replicating data from a source replication server to multiple target replication servers.

Streamline and improve replication by using replication storage rules

IBM Spectrum® Protect 8.1.13 introduces a feature for replicating data by defining replication storage rules and associated subrules. This feature streamlines the configuration process and supports fine-tuning of replication operations. In addition, the feature supports enhanced protection of data in directory-container storage pools.

Updated storage references based on the next generation IBM FlashSystem 5000 and IBM FlashSystem 5200

The small, medium, and large storage references have been updated to the IBM FlashSystem 5015, 5035, and 5200 respectively which provide a scalable hybrid storage solution that is fast, affordable, and reliable.

Updated storage references based on the IBM Elastic Storage® System 5000 and 3200

The powerful IBM Elastic Storage System models 5000 and 3200 are combined to simplify bringing the power and flexibility of the IBM Spectrum Scale file system for running multiple IBM Spectrum Protect servers across a global namespace. This update provides storage for both database and directory container pool storage.

The blueprint configuration automation no longer provides default passwords

The default values for passwords set during server configuration have been eliminated to improve security. The user is now required to provide values for all passwords set during the initial configuration. In addition, the blueprint configuration automation now enforces the new minimum password length of 15 characters introduced in IBM Spectrum Protect 8.1.16.

Technical and other updates were made throughout the book. Look for the vertical bar (|) in the margin.

Chapter 1. Introduction

This document provides detailed steps to build a extra small, small, medium, or large IBM Spectrum Protect server with disk-only storage that uses data deduplication on a Linux® x86 system.

Two options for the storage architecture are included:

- IBM FlashSystem® with Fibre Channel attachments
- IBM Elastic Storage System with an Ethernet or Infiniband attachment.

By following prerequisite steps precisely, you can set up hardware and prepare your system to run the IBM Spectrum Protect Blueprint configuration script, `TSMserverconfig.pl`, for a successful deployment. The settings and options that are defined by the script are designed to ensure optimal performance, based on the size of your system.

Overview

The following roadmap lists the main tasks that you must complete to deploy a server:

1. Determine the size of the configuration that you want to implement.
2. Review the requirements and prerequisites for the server system.
3. Set up the hardware by using detailed blueprint specifications for system layout.
4. Configure the hardware and install the Red Hat Enterprise Linux x86-64 operating system.
5. Prepare storage for IBM Spectrum Protect.
6. Run the IBM Spectrum Protect workload simulation tool to verify that your configuration is functioning properly.
7. Install the IBM Spectrum Protect backup-archive client.
8. Install a licensed version of the IBM Spectrum Protect server.
9. Run the Blueprint configuration script to validate your hardware configuration, and then configure the server.
10. Complete post-configuration steps to begin managing and monitoring your server environment.

Deviating from the Blueprints

Avoid deviations from the Blueprints. Deviations can result in poor performance or other operational problems. Some customization, including substituting comparable server and storage models from other manufacturers, can be implemented, but care must be taken to use components with equivalent or better performance. Avoid the following deviations:

- Running multiple IBM Spectrum Protect server instances on the same operating system on the same computer.
- Reducing the number of drives by substituting larger capacity drives.
- Using the capacity-saving features of storage systems including thin provisioning, compression, or data deduplication. These features are provided by the IBM Spectrum Protect software and redundant use in the storage system can lead to performance problems.

Chapter 2. Implementation requirements

Select the appropriate size for your IBM Spectrum Protect environment and then review requirements for hardware and software.

Use [Table 1 on page 3](#) to select the server size, based on the amount of data that you manage. Both the total managed data and daily amount of new data are measured before data deduplication.

Data amounts in the table are based on the use of directory-container storage pools with inline data deduplication, a feature that was introduced in IBM Spectrum Protect Version 7.1.3. The blueprints are also designed to use inline storage pool compression, a feature that was introduced in IBM Spectrum Protect V7.1.5.

Tip: Before you configure a solution, learn about container storage pools. See [Directory-container storage pools FAQs](#).

If your total managed data is in this range	And the amount of new data that you back up with one replication copy is in this range	The amount of new data that you back up with two replication copies is in this range	Build a server of this size
10 TB - 40 TB	Up to 1 TB per day	Up to 0.6 TB per day	Extra Small
60 TB - 240 TB	Up to 10 TB per day	Up to 6 TB per day	Small
360 TB - 1440 TB	10 - 30 TB per day	6 - 18 TB per day	Medium
1000 TB - 4000 TB	30 - 100 TB per day	18 - 60 TB per day	Large

The *daily ingestion rate* is the amount of data that you back up each day. The daily ingestion needs to be completed in a backup window that leaves enough time remaining in the day to complete maintenance tasks. For optimum performance, split the tasks of backing up and archiving client data, and performing server data maintenance into separate time windows. The daily ingestion amounts in [Table 1 on page 3](#) are based on test results with 128 MB sized objects, which are used by IBM Spectrum Protect for Virtual Environments assuming a backup window of eight hours. The daily ingestion amount is stated as a range because backup throughput, and the time that is required to complete maintenance tasks, vary based on workload.

If a server is used to both accept backup data, and receive replicated data from other servers, more planning is needed. Any data that is received through replication must be considered as part of the daily backup amount. For example, a server that receives 25 TB of new backup data and 15 TB of new replication data daily has a total ingestion rate of 40 TB per day. Optionally, backup data and data received through replication can be placed in separate directory container storage pools.

Remember: If you are planning to create two replication copies of the backup data, you will need to consider it while selecting the size of the server. The daily amount of backup data has to be decreased to reduce the amount of time required to back up data. This is done to compensate for the additional time needed to create the second replication copy.

Not every workload can achieve the maximum amount in the range for daily backups. The range is a continuum, and placement within the range depends on several factors:

Major factors

- Average object size. Workloads with smaller average object sizes, such as those that are common with file server backups, typically have smaller backup throughputs. If the average object size is less than 128 KB, daily backup amounts are likely to fall in the lower 25% of the range. If the average object size is larger, for example, 512 KB or more, backup throughputs are greater.

- Daily data reduction. When data is reduced by using data deduplication and compression, less data must be written to storage pools. As a result, the server can handle larger amounts of daily data ingestion.

Additional factors

- Data deduplication location. By using client-side data deduplication, you reduce the processing workload on the server. As a result, you can increase the total amount of data that is deduplicated daily.
- Network performance. By using efficient networks, you can back up and replicate more data daily.

Additionally, including optional features in the solution, such as making a copy of the container storage pool to tape storage, will require adjustments to the maximum amount of new backup data that can be processed per day. The amount of time required to complete the optional data copy or movement activities needs to be considered in evaluating the daily ingest limit for the server.

To better understand the factors that affect the maximum amount of daily data ingestion, review the following figure:




Figure 1. Range for daily data ingestion in a large system

Total managed data is the amount of data that is protected. This amount includes all versions. A range is provided because data processing responds differently to data deduplication and compression, depending on the type of data that is backed up. The smaller number in the range represents the physical capacity of the IBM Spectrum Protect storage pool. Although the use of inline compression does not result in additional growth of the IBM Spectrum Protect database, compression might result in the ability to store more data in the same amount of storage pool space. In this way, the amount of total managed data can increase causing more database space to be used.

To estimate the total managed data for your environment, you must have the following information:

- The amount of client data (the front-end data amount) that will be protected
- The number of days that backup data must be retained
- An estimate of the daily change percentage
- The backup model that is used for a client type, for example, incremental-forever, full daily, or full periodic

If you are unsure of your workload characteristics, use the middle of the range for planning purposes.

You can calculate the total managed data for different types of clients in groups and then add the group results.

Client types with incremental-forever backup operations

Use the following formula to estimate the total managed data:

```
Frontend + (Frontend * changerate * (retention - 1))
```

For example, if you back up 100 TB of front-end data, use a 30-day retention period, and have a 5% change rate, calculate your total managed data as shown:

```
100 TB + (100TB * 0.05 * (30-1)) = 245 TB total managed data
```

Client types with full daily backup operations

Use the following formula to estimate the total managed data:

```
Frontend * retention * (1 + changerate)
```

For example, if you back up 10 TB of front-end data, use a 30-day retention period, and have a 3% change rate, calculate your total managed data as shown:

```
10 TB * 30 * (1 + .03) = 309 TB total managed data
```

To efficiently maintain periodic copies of your data to meet long-term retention requirements, you can use the retention set feature. Retention sets are created from existing backups without requiring data to be redundantly sent to the IBM Spectrum Protect server. Retention sets can either be created in-place by maintaining the existing backups for multiple retention requirements, or with copies made to tape media. In-place retention sets will increase the amount of total managed data requiring additional storage pool and database space. Retention set copies will require space in a retention pool, but have a very minimal impact to database space.

Hardware and software prerequisites

Before you set up your system and configure IBM Spectrum Protect, you must meet hardware and software prerequisites.

The hardware requirements which follow specify server and storage specifications to meet different sized workloads. References to CPU core requirements are referring to physical CPU cores, and not virtual CPU threads.

Hardware requirements

You must acquire hardware that is based on scale size requirements. You can choose equivalent or better components than what are listed.

The following topics list the hardware requirements for a extra small, small, medium, or large configuration. The tables contain detailed descriptions, including part numbers and quantities for IBM® components that are used in the storage configuration blueprints.

The system memory sizes that are provided are recommendations for optimal performance. They are not minimum requirements. Memory recommendations account for using both data deduplication and node replication with a database that is near maximum capacity. Some workloads can tolerate smaller amounts of memory. When node replication is not used, the memory recommendations can be reduced by 25%.

The hardware specifications that are provided are current at the time of publishing. Part substitutions might be required, depending on hardware availability over time. Be cautious if you plan to substitute a smaller quantity of larger drives, particularly for the database. A smaller quantity of drives might not provide comparable performance.

IBM FlashSystem storage systems are designed to provide a simple, high-performance solution for managing block-based storage. For this reason, FlashSystem storage systems are suited for use by the IBM Spectrum Protect server for both the database and storage pool. For more information about FlashSystem features, see [IBM Flash Storage family](#).

Note: The IBM FlashSystem 92 drive expansion racks require more rack depth than other disk expansion options. Review the product specifications for rack requirements to make sure racks that support the required depth are available.

Recent IBM Spectrum Protect releases introduce new capabilities for moving or copying data to tape storage. If you are planning to include the optional features of tiering to tape, making retention set copies to tape, or copying the container pool to tape, you will need to increase the number of fibre channel ports in the configuration. Consider the following:

- Fibre channel traffic for disk access and tape access should be zoned to different fibre channel ports rather than sharing ports.
- The data being moved or copied to tape is reconstructed and uncompressed to its original size before being transferred to tape. For this reason, more fibre channel port capacity is needed for the tape access than the disk access. For a system which requires a single port for disk access, at least two additional ports dedicated for tape access will be required. For a system which requires two ports for disk access, at least four additional ports are required for tape access.

The tables in the following topics have abbreviated part lists, which include only *major* components. Work with your hardware supplier to ensure that your configuration is complete.

Hardware requirements for extra small systems

Extra small systems can be deployed as virtual machines which meet the specifications for virtual hardware listed below.

Extra small systems have been tested as VMware virtual machines.

Hardware	Requirements	Blueprint component	Detailed description
Server and network	<ul style="list-style-type: none"> • Four virtual processor cores, 1.7 GHz or faster • 24 GB RAM • 1Gb or 10 Gb Ethernet 	VMware ESXi Version 6.7 or 7.0	Virtual machine with a virtual hardware level of 13 or newer. VMware tools must be installed.
			4-core virtual CPU
			24 GB virtual RAM
			Virtual network adapter of type E1000E
Disks for storage	Virtual disks can either be assigned as RDM disks or as virtual disks. Virtual disks must be thickly provisioned, and VMware snapshots should not be taken of the virtual disks.	When using virtual disks, create the virtual disks for the operating system, database, and storage pools in different VMware datastores.	Operating system disk <ul style="list-style-type: none"> • Size: 90 GB • Qty: 1
			Database <ul style="list-style-type: none"> • Size: 100 GB • Qty: 2
			Active log <ul style="list-style-type: none"> • Size: 30 GB • Qty: 1
			Archive log <ul style="list-style-type: none"> • Size: 250 GB • Qty: 1
			Database backup <ul style="list-style-type: none"> • Size: 1000 GB • Qty: 1
			Storage pool <ul style="list-style-type: none"> • Size: 5000 GB • Qty: 2

Hardware requirements for small systems

You must acquire hardware that is based on scale size requirements. You can choose equivalent or better components than what are listed.

Server references are provided using Lenovo ThinkSystem SR650 servers. Equivalent x86_64 servers from other manufactures can be substituted. :

- For Lenovo product information, see [Lenovo ThinkSystem SR650 Rack Server](#). For hardware requirements, see [Table 3 on page 7](#).

Table 3. Hardware requirements for a small system that uses a Lenovo server

Hardware	Requirements	Blueprint component	Detailed description	Quantity	Part number
Server and network	<ul style="list-style-type: none"> • 16 processor cores, 1.7 GHz or faster • 64 GB RAM • 10 Gb Ethernet • 8 Gb or 16 Gb Fibre Channel adapter 	Lenovo ThinkSystem SR650	Lenovo ThinkSystem SR650	1	7X06CT01W
			8-core Intel Xeon Bronze 3206 1.9 GHz processor	2	B7N3
			8 GB TruDDR4 2933 MHz memory	8	B4H1
			Mellanox Connect X-4 L 10/25GbE 2-port PCIe Ethernet adapter	1	AUAJ
			Emulex 16 Gb Gen6 FC dual-port HBA	1	ATZV
			RAID 530-16i PCIe 12 Gb adapter	1	B6DJ
			300 GB 10K SAS 12 Gb Hot Swap 512n HDD	2	B1JJ
Disks for storage	<ul style="list-style-type: none"> • 16 Gb host interface • Database and active log disks: 800 GB SSD flash drives • Storage pool disks: 4 TB NL-SAS 	IBM FlashSystem 5015	IBM FlashSystem 5015 SFF Control	1	2072-2N4
			16 Gb Fibre Channel adapter pair	1	ALBB
			V5000E CACHE UPGRADE	1	ALGA
			800 GB 3DWPD 2.5 flash drive	4	AL8A
			IBM FlashSystem 5000 Large form-factor (LFF) Expansion Enclosure	2	2072-12G
			0.6 m SAS Cable (mSAS HD)	4	ACUA
			4 TB 7.2 K 3.5-inch NL HDD	24	AL39

Hardware requirements for medium systems

You must acquire hardware that is based on scale size requirements. You can choose equivalent or better components than what are listed.

Server references are provided using Lenovo ThinkSystem SR650 servers. Equivalent x86_64 servers from other manufactures can be substituted.

- For Lenovo product information, see [Lenovo ThinkSystem SR650 Rack Server](#). For hardware requirements, see [Table 4 on page 8](#).

Table 4. Hardware requirements for a medium system that uses a Lenovo server

Hardware	Requirements	Blueprint component	Detailed description	Quantity	Part number
Server and network	<ul style="list-style-type: none"> • 20 processor cores, 2.2 GHz or faster • 192 GB RAM • 10 Gb Ethernet • 8 Gb or 16 Gb Fibre Channel adapter 	Lenovo ThinkSystem SR650	Lenovo ThinkSystem SR650	1	7X06CTO1W
			10-core Intel Xeon Silver 4210 2.2 GHz processor	2	B4HS
			16 GB TruDDR4 2933 MHz memory	12	AUNC
			Mellanox Connect X-4 L 10/25GbE 2-port PCIe Ethernet adapter	1	AUAJ
			Emulex 16 Gb Gen6 FC dual-port HBA	2	ATZV
			RAID 730-8i PCIe 12 Gb adapter	1	B4RQ
			300 GB 10K SAS 12 Gb Hot Swap 512n HDD	2	AULY
Disks for storage	<ul style="list-style-type: none"> • 16 Gb host interface • Database and active log disks: 1.92 TB SSD • Storage pool, archive log, and database backup disks: 6 TB NL-SAS 	IBM FlashSystem 5035	IBM FlashSystem 5035 SFF Control	1	2072-3N4
			16 GB Fibre Channel adapter pair	1	ALBB
			V5000E CACHE UPGRADE	1	ALGA
			1.92 TB 2.5-inch flash drive	6	AL80
			5000 HD large form-factor (LFF) expansion	1	2072-92G
			6 TB 7.2 K 3.5-inch NL HDD	92	AL47
			3 m 12 Gb SAS cable (mSAS HD)	2	ACUC

Hardware requirements for large systems

You must acquire hardware that is based on scale size requirements. You can choose equivalent or better components than what are listed.

Server references are provided using Lenovo ThinkSystem SR650 servers. Equivalent x86_64 servers from other manufactures can be substituted.

- For Lenovo product information, see [Lenovo ThinkSystem SR650](#). For hardware requirements, see [Table 5 on page 8](#).

Table 5. Hardware requirements for a large system that uses a Lenovo server

Hardware	Requirements	Blueprint component	Detailed description	Quantity	Part number
Server and network	<ul style="list-style-type: none"> • 44 processor cores, 2.1 GHz or faster • 576 GB RAM • 10 or 25 Gb Ethernet • 16 Gb Fibre Channel adapter 	Lenovo ThinkSystem SR650	Lenovo ThinkSystem SR650	1	7X06CTO1W
			22-core Intel Xeon Gold 6238 2.1 GHz processor	2	B6CJ
			32 GB TruDDR4 2933 MHz Memory and 16GB TruDDR4 2933 MHz Memory	12 each	B4H3 and B4H2
			Mellanox Connect X-4 L 10/25GbE 2-port PCIe Ethernet adapter	2	AUAJ
			Emulex 16 Gb Gen6 FC dual-port HBA	2	ATZV
			RAID 730-8i PCIe 12 Gb adapter	1	B4RQ
			300 GB 10K SAS 12 Gb Hot Swap 512n HDD	3 ¹	AULY

Table 5. Hardware requirements for a large system that uses a Lenovo server (continued)

Hardware	Requirements	Blueprint component	Detailed description	Quantity	Part number
Disks for storage	<ul style="list-style-type: none"> Database and active log disks: 1.92 TB NVME FLASH DRIVE Storage pool, archive log, and database backup disks: 8 TB NL-SAS drives 	IBM FlashSystem 5200	IBM FlashSystem 5200 NVMe Control Enclosure	1	4662-6H2
			16 Gb FC 4 Port Adapter Cards (pair)	1	ALBJ
			IBM 512 GB Base Cache	1	ALG1
			1.92 TB 2.5-inch NVME flash drive	9	AGT2
			IBM FlashSystem 5200 High Density Expansion Enclosure	2	4662-92G
			8 TB 7,200 rpm 12 Gb SAS NL 3.5 Inch HDD	184	AL48
			3M 12Gb SAS CABLE MSAS HD	4	ACUC

1. Two of the three 300 GB internal hard disks are configured in a RAID 1 pair, and the third drive is assigned as a spare. If a spare is not needed based on business requirements, the system can be configured with only two drives.

Hardware requirements for IBM Elastic Storage System

Review the hardware requirements. You can choose equivalent or better components than what are listed.

The high-level components that are needed to build a large system with IBM Elastic Storage System storage are listed in the following table. To complete the configuration tasks by using this information, contact your IBM representative or Business Partner.

The following reference shows an IBM Elastic Storage System 5000 SL6 model that uses a 10 TB drive size. This configuration provides capacity that is sufficient for four or more IBM Spectrum Protect servers. Other configurations with lower performance levels, smaller capacity, or both are possible by using other models with different expansion counts and drive sizes.

Large system

Table 6. Hardware requirements for a large system that uses IBM Elastic Storage System

Hardware	Requirements	Blueprint component	Detailed description	Quantity	Part number
Server and network	<ul style="list-style-type: none"> 44 processor cores, 2.1 GHz or faster 576 GB RAM 100 Gb Infiniband for storage IO 4 x 25GbE for client backups and replication 	Lenovo ThinkSystem SR650	Lenovo ThinkSystem SR650	1	7X06CTO1W
			22-core Intel Xeon Gold 6238 2.1 GHz processor	2	B6CJ
			32 GB TruDDR4 2933 MHz Memory and 16GB TruDDR4 2933 MHz Memory	12 each	B4H3 and B4H2
			Mellanox ConnectX-4 2x100GbE/EDR IB QSFP28 VPI Adapter	1	
			Mellanox Connect X-4 L 10/25GbE 2-port PCIe Ethernet adapter	2	AUAJ
			RAID 730-8i PCIe 12 Gb adapter	1	B4RQ
			300 GB 10K SAS 12 Gb Hot Swap 512n HDD	3 ²	AULY

Table 6. Hardware requirements for a large system that uses IBM Elastic Storage System (continued)

Hardware	Requirements	Blueprint component	Detailed description	Quantity	Part number
Storage system	<ul style="list-style-type: none"> Storage pool disks: 10TB Enterprise HDD Database disks: 3.84 TB PVMme Flash Drive 	IBM Elastic Storage System model 5000 SL6 and model 3200	IBM Elastic Storage System system	1	
			Data server	2	5105-22E
			Management server	1	5105-22E
			Storage Expansion	6	5147-092
			10 TB Enterprise HDD	550	AJNX
			800 GB 6Gb SAS SSD	2	AJNV
			IBM ESS 3200	1	5141-FN1
			3.84 TB PCIe Gen4 NVMe Flash Drive	24	AJP4

1. There are other networking configurations that are possible.
2. Two of the three 300 GB internal hard disks are configured in a RAID 1 pair, and the third drive is assigned as a spare. If a spare is not needed based on business requirements, the system can be configured with only two drives.

Software requirements

You must install the Linux operating system and the IBM Spectrum Protect server and backup-archive client.

The following versions are required:

- Red Hat Enterprise Linux x86_64, Version 7.8 or later or Red Hat Enterprise Linux x86_64, Version 8.5 or later.
- IBM Spectrum Protect V8.1.12 or later backup-archive client.
- A licensed version of IBM Spectrum Protect is required to run the Blueprint configuration script. To obtain critical fixes, install IBM Spectrum Protect V8.1.14.100 or later. RHEL 8 support is available starting with IBM Spectrum Protect was V8.1.11. At the time of publication, the latest level of IBM Spectrum Protect was V8.1.16.
- The Blueprint configuration script V5.1 or later.

Planning worksheets

Use the planning worksheets to record values that you use when you complete the steps to set up your system and then configure the IBM Spectrum Protect server. The preferred method is to use the default values that are listed in the worksheets.

Default values in the following tables correspond to the default values that are used by the Blueprint configuration script to configure the server. By using these values to create your file systems and directories, you can accept all defaults for the configuration when you run the script. If you create directories or plan to use values that do not match the defaults, you must manually enter those values for the configuration.

Planning worksheets for IBM FlashSystem configurations

Use Table 7 on page 11 to plan for the file systems and directories that you create during system setup. All directories that you create for the server must be empty.

Table 7. Values needed for preconfiguration

Item	Default value	Your value	Directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	This value is used when you install and configure the operating system and is assigned by the Blueprint configuration script during configuration of the server. If you want to use a value other than the default, you can specify a number in the range 1024 - 32767.
Directory for the server instance	/home/tsminst1/tsminst1		100 GB	If you change the value for the server instance directory from the default, modify the IBM Db2® instance owner ID in Table 8 on page 12 as well.
Directory for server installation	/		30 GB	The directory size value is the minimum available space that you must have. For more information about system requirements on the Linux x86_64 operating system, see technote 1204361 .
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Extra small: 30 GB • Small and medium: 140 GB • Large: 550 GB 	
Directory for the archive log	/tsminst1/TSMarchlog		<ul style="list-style-type: none"> • Extra small: 250 GB • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	
Directories for the database	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 and so on.		Minimum total space for all directories: <ul style="list-style-type: none"> • Extra small: At least 200 GB • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB 	Create a minimum number of file systems for the database, depending on the size of your system: <ul style="list-style-type: none"> • Extra small: At least 1 file system • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 and so on.		Minimum total space for all directories: <ul style="list-style-type: none"> • Extra small: At least 10 TB • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Extra small: At least 2 file systems • Small: At least 2 file systems • Medium: At least 10 file systems • Large: At least 30 file systems

Table 7. Values needed for preconfiguration (continued)

Item	Default value	Your value	Directory size	Notes
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03 and so on.		Minimum total space for all directories: <ul style="list-style-type: none"> • Extra small: At least 1 TB • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB 	Create a minimum number of file systems for backing up the database, depending on the size of your system: <ul style="list-style-type: none"> • Extra small: At least 1 file system • Small: At least 2 file systems • Medium: At least 3 file systems • Large: At least 3 file systems The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.

Use Table 8 on page 12 when you run the Blueprint configuration script to configure the server. The preferred method is to use the default values, except where noted.

Table 8. Values needed for the server configuration

Item	Default value	Your value	Notes
Db2 instance owner ID	tsminst1		If you changed the value for the server instance directory in Table 7 on page 11 from the default, modify the value for the Db2 instance owner ID as well.
Db2 instance owner password	There is no default for this value.		The user is required to select a value for the instance owner password. Ensure that you record this value in a secure location.
Primary group for the Db2 instance owner ID	tsmsivrs		
Server name	The default value for the server name is the system host name.		
Server password	There is no default for this value.		The user is required to select a value for the server password. Ensure that you record this value in a secure location.
Administrator ID (user ID for the server instance)	admin		
Administrator ID password	There is no default for this value.		The user is required to select a value for the administrator password. Ensure that you record this value in a secure location.
Schedule start time	22:00		The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. These operations are usually completed during the nightly schedule window. Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.

Planning worksheets for server configurations

Item	Default value	Your value	Directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	This value is used when you install and configure the operating system and is assigned by the Blueprint configuration script during configuration of the server. If you want to use a value other than the default, you can specify a number in the range 1024 - 32767.
Directory for the server instance	/esstsm1/tsminst1/tsminst1		Not applicable	If you change the value for the server instance directory from the default, modify the Db2 instance owner ID in Table 8 on page 12 as well.
Directory for the active log	/esstsmflash/tsminst1/database/alog		Not applicable	
Directory for the archive log	/esstsm1/tsminst1/database/archlog		Not applicable	
Directories for the database	/esstsmflash/tsminst1/database/ TSMdbspace00 /esstsmflash/tsminst1/database/ TSMdbspace01 /esstsmflash/tsminst1/database/ TSMdbspace02 and so on.		<ul style="list-style-type: none"> • Medium: 2 TB • Large: 4 TB 	Create at least 8 directories for the server database.
Directories for storage	/esstsm1/tsminst1/deduppool/		Not applicable	
Directories for database backup	/esstsm1/tsminst1/dbback		Not applicable	

Item	Default value	Your value	Notes
Db2 instance owner ID	tsminst1		If you changed the value for the server instance directory in Table 7 on page 11 from the default, modify the value for the Db2 instance owner ID as well.
Db2 instance owner password	There is no default for this value.		The user is required to select a value for the instance owner password. Ensure that you record this value in a secure location.
Primary group for the Db2 instance owner ID	tsmsrvrs		

Table 10. Values needed for the server configuration (continued)

Item	Default value	Your value	Notes
Server name	The default value for the server name is the system host name.		
Server password	There is no default for this value.		The user is required to select a value for the server password. Ensure that you record this value in a secure location.
Administrator ID (user ID for the server instance)	admin		
Administrator ID password	There is no default for this value.		The user is required to select a value for the administrator password. Ensure that you record this value in a secure location.
Schedule start time	22:00		<p>The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. These operations are usually completed during the nightly schedule window.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Chapter 3. Storage configuration blueprints

After you acquire hardware for the scale of server that you want to build, you must prepare your storage to be used with IBM Spectrum Protect. Configuration blueprints provide detailed specifications for storage layout. Use them as a map when you set up and configure your hardware.

Specifications in [“Hardware requirements” on page 5](#) and the default values in the [“Planning worksheets” on page 10](#) were used to construct the blueprints for small, medium, and large systems. If you deviate from those specifications, you must account for any changes when you configure your storage.

Note: The IBM FlashSystem configurations implement fully-allocated volumes that do not use hardware data reduction techniques including compression and deduplication. The IBM Spectrum Protect software will perform the data reduction, and redundantly performing these tasks in the storage system will result in performance problems.

If you are configuring a system with IBM Elastic Storage System, see [“IBM Elastic Storage System” on page 21](#).

Distributed arrays

You can use the distributed arrays feature with NL-SAS drives to achieve faster drive rebuild times in case of a disk failure. FlashSystem distributed arrays, which contain 4 - 128 drives, also contain rebuild areas that are used to maintain redundancy after a drive fails. The distributed configuration can reduce rebuild times and decrease the exposure of volumes to the extra workload of recovering redundancy. If you plan to use the 92-drive FlashSystem expansions, the preferred method is to create two 46-drive distributed RAID 6 arrays per expansion.

If you are using a disk system that does not support distributed arrays, you must use traditional storage arrays. For instructions about configuring traditional storage arrays, see the *Blueprint and Server Automated Configuration, Version 2 Release 3* guide for your operating system at the [IBM Spectrum Protect Blueprints website](#).

Tip: Earlier versions of the blueprints are available at the bottom of the blueprint web page.

FlashSystem layout requirements

A *managed disk*, or *MDisk*, is a logical unit of physical storage. In the blueprint configurations, MDisk are internal-storage RAID arrays and consist of multiple physical disks that are presented as logical volumes to the system. When you configure the disk system, you will create MDisk groups, or data storage pools, and then create MDisk arrays in the groups.

The medium and large blueprint configurations include more than one MDisk distributed array and combine the MDisks together into a single MDisk group or storage pool. In previous blueprint versions, a one-to-one mapping exists between MDisks and MDisk groups. Sharing a common storage pool for multiple arrays is not required for disk systems which do not support this or for configurations that were implemented to the earlier blueprint design.

Volumes, or LUNs, belong to one MDisk group and one I/O group. The MDisk group defines which MDisks provide the storage that makes up the volume. The I/O group defines which nodes provide I/O access to the volume. When you create volumes, make them fully allocated with a vdev type of striped. For IBM FlashSystem hardware, select the generic volume type when you create volumes.

[Table 11 on page 15](#) and [Table 12 on page 16](#) describe the layout requirements for MDisk and volume configuration in the storage blueprints.

Component	Details
Server storage requirement	How the storage is used by the IBM Spectrum Protect server.

Table 11. Components of MDisk configuration (continued)

Component	Details
Disk type	Size and speed for the disk type that is used for the storage requirement.
Disk quantity	Number of each disk type that is needed for the storage requirement.
Hot spare coverage	Number of disks that are reserved as spares to take over in case of disk failure. For distributed arrays this represents the number of rebuild areas.
RAID type	Type of RAID array that is used for logical storage.
RAID array quantity and DDM per array	Number of RAID arrays to be created, and how many disk drive modules (DDMs) are to be used in each of the arrays.
Usable size	Size that is available for data storage after accounting for space that is lost to RAID array redundancy.
Suggested MDisk names	Preferred name to use for MDisks and MDisk groups.
Usage	IBM Spectrum Protect server component that uses part of the physical disk.

Table 12. Components of volume (LUN) configuration

Component	Details
Server storage requirement	Requirement for which the physical disk is used.
Volume name	Unique name that is given to a specific volume.
Quantity	Number of volumes to create for a specific requirement. Use the same naming standard for each volume that is created for the same requirement.
Uses MDisk group	The name of the MDisk group from which the space is obtained to create the volume.
Size	The size of each volume.
Intended server mount point	The directory on the IBM Spectrum Protect server system where the volume is mounted. If you plan to use directories other than the defaults that are configured by the Blueprint configuration script, you must also use those directory values when you configure your hardware. In this case, do not use the values that are specified in the blueprints.
Usage	IBM Spectrum Protect server component that uses part of the physical disk.

FlashSystem volume protection feature

The IBM FlashSystem volume protection feature is a safeguard that prevents unintended deletion of volumes containing important data when there has been recent I/O against the volumes. Activate this feature to protect the volumes used with IBM Spectrum Protect. The volume protection feature is not on by default, and must be enabled for each storage pool from the IBM FlashSystem user interface.

Small FlashSystem configuration

A small-scale system is based on IBM FlashSystem 5015 storage. One dual control enclosure and two expansion enclosures contain IBM Spectrum Protect data.

Logical layout

Figure 2 on page 17 shows the small system layout and how server and storage hardware is connected to clients. A single cluster and I/O group are used in this configuration. The small system configuration was tested with 8 Gb Fibre Channel connections made directly from the host to the FlashSystem 5015 system without a SAN switch. The following image depicts a configuration that uses a Lenovo ThinkSystem SR650 server.

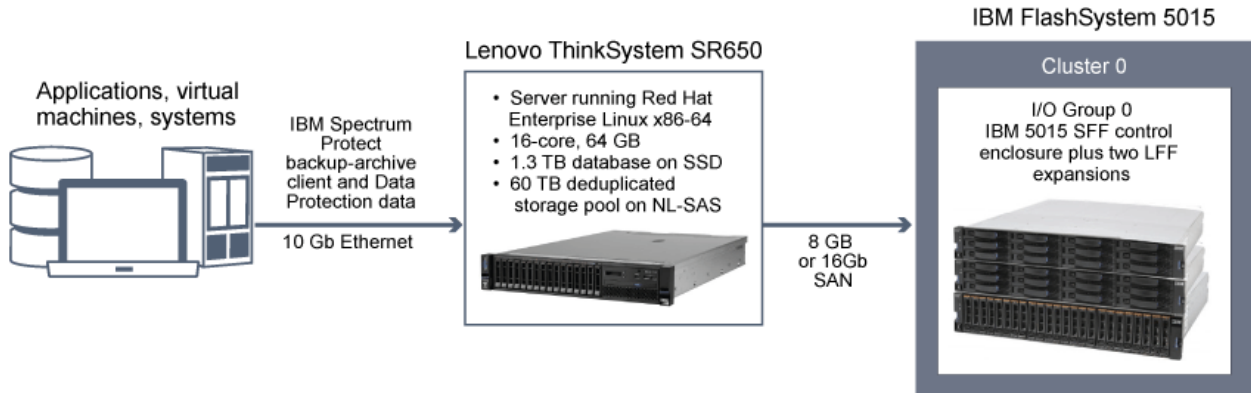


Figure 2. Logical layout for a small system

Storage configuration

Table 13 on page 17 and Table 14 on page 18 show the detailed layout for each IBM Spectrum Protect storage requirement on a small system.

Table 13. MDisk configuration

Server storage requirement	Disk type	Disk quantity	Hot spare coverage	RAID type	RAID array quantity	Usable size	Suggested MDisk group and array names	Usage
Database	800 GB SSD	4	1 rebuild areas=1	DRAID 5 ¹	1 4 DDM	1.45 TB	db_grp0 db_array0	Database container
Storage pool	4 TB 7.2k rpm NL-SAS HDD	24	rebuild areas=1	DRAID 6 ²	1 24 DDM	67 TB	stgpool_grp0 stgpool_array0	Storage pool

1. Distributed RAID 5, stripewidth=3, rebuild areas=1.
2. Distributed RAID 6, stripewidth=12, rebuild areas=1.

Table 14. Fully allocated volume configuration

Server storage requirement	Volume name	Quantity	Uses MDisk group	Size	Intended server mount point	Usage
Database	db_00 - db_03	4	db_grp0	335.25 GB each	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03	Database container
	alog	1	db_grp0	145.25 GB	/tsminst1/TSMalog	Active log
	archlog	1	stgpool_grp0	1.19 TB	/tsminst1/TSMarchlog	Archive log
	backup_0 - backup_1	2	stgpool_grp0	3.15 TB each	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01	Database backup
Storage pool	filepool_00 - filepool_03	4	stgpool_grp0	15.12 TB each	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03	IBM Spectrum Protect file systems for a directory-container storage pool

Medium FlashSystem configuration

A medium-scale system is based on IBM FlashSystem 5035 hardware. One dual control enclosure and one large capacity expansion enclosure contain IBM Spectrum Protect data.

Logical layout

Figure 3 on page 18 shows the medium system layout and how server and storage hardware is connected to clients. A single cluster and I/O group are used. The medium system configuration was tested by using a SAN switch with 16 Gb Fibre Channel connections and two bonded 10 Gb Ethernet connections. The image depicts a configuration that uses a Lenovo ThinkSystem SR650 server.

The tables show multiple distributed arrays that are members of the same FlashSystem storage pool. Alternatively, you can create split the arrays into separate storage pools.

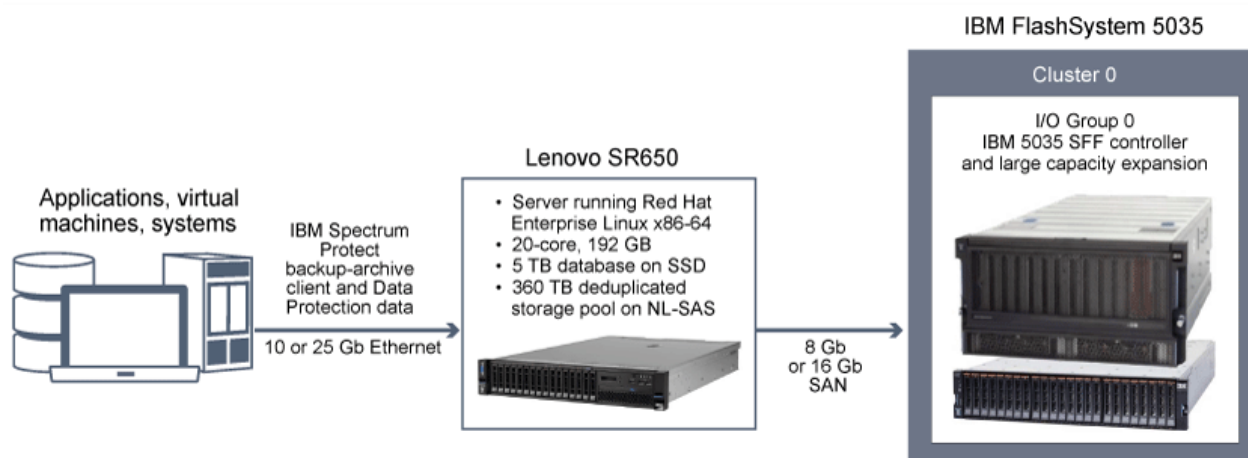


Figure 3. Logical layout for a medium system

Storage configuration

Table 15 on page 19 and Table 16 on page 19 show the detailed layouts for MDisk and volume configurations on a medium system. The following array configuration requires the default FlashSystem memory allocation for RAID to be increased, as described in Step “2” on page 68.

Server storage requirement	Disk type	Disk quantity	Hot spare coverage	RAID type	RAID array quantity	Usable size	Suggested MDisk group and array names	Usage
Database	1.92 TB SSD	6	1 rebuild-areas = 1	DRAID6 ¹	1 6 DDM	5.16 TB	db_grp0 db_array0	Database and active log
Storage pool	6 TB NL-SAS	92	4 rebuild-areas = 2	DRAID6 ²	2 46 DDM each	197.91 TB each	stgpool_grp0 stgpool_array0 and stgpool_array1	Storage pool, archive log, and database backups

1. Distributed RAID 6, stripe width=5, rebuild areas=1.
2. Distributed RAID 6, stripe width=12, rebuild areas=2.

Server storage requirement	Volume name	Quantity	Uses MDisk group	Size	Intended server mount point	Usage
Database	db_00 db_01 db_02 db_03 db_04 db_05 db_06 db_07	8	db_grp0	642.1 GB each	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 /tsminst1/TSMdbspace04 /tsminst1/TSMdbspace05 /tsminst1/TSMdbspace06 /tsminst1/TSMdbspace07	Database
	alog	1	db_grp0	147 GB	/tsminst1/TSMalog	Active log
	archlog_00	1	stgpool_grp0	2 TB	/tsminst1/TSMarchlog	Archive log
	backup_00 backup_01 backup_02	3	stgpool_grp0	15 TB each	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02	Database backup
Storage pool	filepool00 - filepool11	12	stgpool_grp0	29.22 TB each	/tsminst1/TSMfile00 /tsminst1/TSMfile01 ... /tsminst1/TSMfile11	IBM Spectrum Protect file systems for a directory-container storage pool

Large FlashSystem configuration

A large-scale system is based on IBM FlashSystem 5200 hardware. One controller with two high-density expansions contains the data.

Logical layout

Figure 4 on page 20 shows the large system layout and how server and storage hardware is connected to clients. Testing for the large system configuration was completed by using a SAN switch with four 16 Gb Fibre Channel connections and four bonded 25 Gb Ethernet connections. The following image depicts a configuration that uses a Lenovo ThinkSystem SR650 server.

The tables show multiple distributed arrays that are members of the same FlashSystem storage pool. Alternatively, you can create split the arrays into separate storage pools.

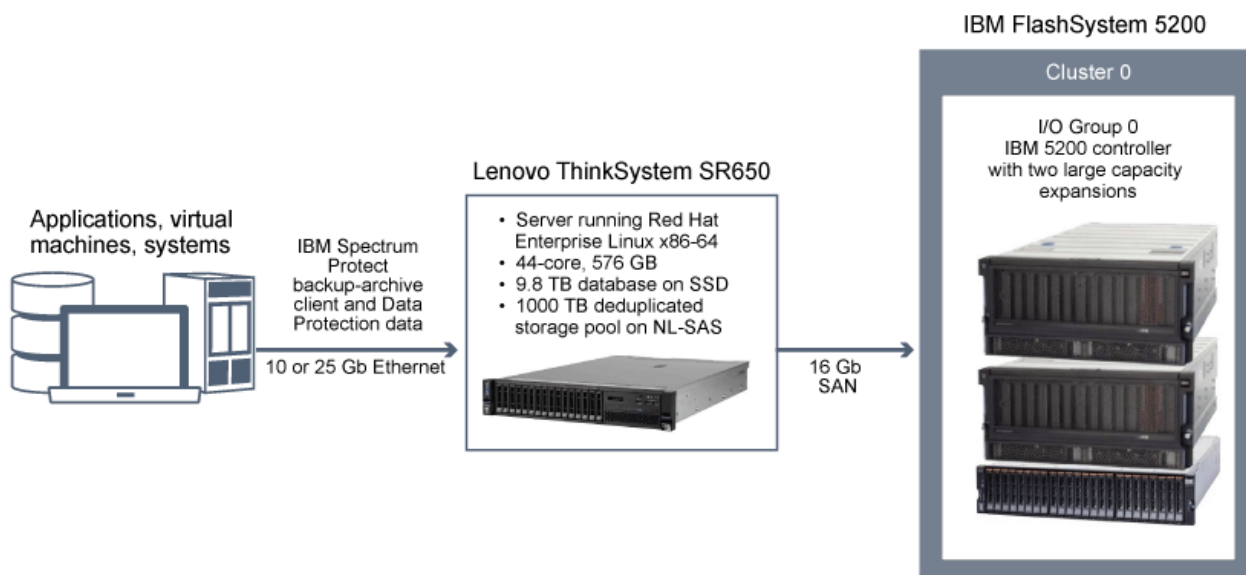


Figure 4. Logical layout for a large system

Storage configuration

Table 17 on page 20 and Table 18 on page 21 show the detailed layouts for MDisk and volume configurations on a large system. To allocate arrays across 184 drives, the memory that is available for RAIDds must be increased to 125 MB, as described in Step “2” on page 68.

Server storage requirement	Disk type	Disk quantity	Hot spare coverage	RAID type	RAID array quantity	Usable size	Suggested MDisk group and array names	Usage
Database	1.92 TB NVME	9	1 rebuild areas = 1	DRAID 6 ¹	1 9 DDM	10.38 TB	db_grp0 db_array0	Database container and active log
Storage pool, archive log, and database backup	8 TB NL-SAS	184	8 rebuild areas = 2 per array	DRAID 6 ²	4 46 DDM each	265.18 TB each	stgpool_grp0 stgpool_array0 stgpool_array1 stgpool_array2 stgpool_array3	Storage pool

1. Distributed RAID 6, stripe width=8, rebuild areas=1.
2. Distributed RAID 6, stripe width=12, rebuild areas=2.

Table 18. Fully allocated volume configuration

Server storage requirement	Volume name	Quantity	Uses MDisk group	Size	Intended server mount point	Usage
Database	db_00 - db_11	12	db_grp0	837.89 GB each	/tsminst1/TSMdbspace00 - /tsminst1/TSMdbspace11	Database
	alog	1	db_grp0	550 GB	/tsminst1/TSMalog	Active log
	archlog	1	stgpool_grp0	4 TB	/tsminst1/TSMarchlog	Archive log
	backup_00 backup_01 backup_02	3	stgpool_grp0	18 TB each	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02	Database backup
Storage pool	filepool_00 - filepool_31	32	stgpool_grp0	31.33 TB each	/tsminst1/TSMfile00 /tsminst1/TSMfile01 ... /tsminst1/TSMfile31	IBM Spectrum Protect file systems for a directory-container storage pool

IBM Elastic Storage System

IBM Elastic Storage System is a scalable storage system that provides a clustered IBM Spectrum Scale file system that is suitable for being shared by several IBM Spectrum Protect servers. By following the instructions, you can configure a large system that uses IBM Elastic Storage System model 5000 SL6 for the directory container storage pool, and the flash storage in the IBM Elastic Storage System model 3200 for database storage.

Storage configuration

Some configuration steps are completed at the factory and by IBM services so that the system will be ready for you to provision storage as single file systems from each disk system to be shared by multiple IBM Spectrum Protect servers. These configuration steps include hardware installation and cabling, software installation on the storage nodes, and configuration of the IBM Elastic Storage System cluster and recovery groups.

Large system layout

Figure 5 on page 22 shows an optional system layout with a single IBM Elastic Storage System providing sufficient storage for four IBM Spectrum Protect servers to share. Each server protects a different set of clients. Clients communicate with an IBM Spectrum Protect server by using 25 Gb Ethernet connections.

The IBM Spectrum Protect servers communicate with the IBM Elastic Storage System system by using 100 Gb Infiniband connections. The IBM Spectrum Scale client is installed on each IBM Spectrum Protect server. High-speed access to the storage is possible by using the network shared disk (NSD) protocol that provides direct access to the IBM Elastic Storage System system through the Ethernet network.

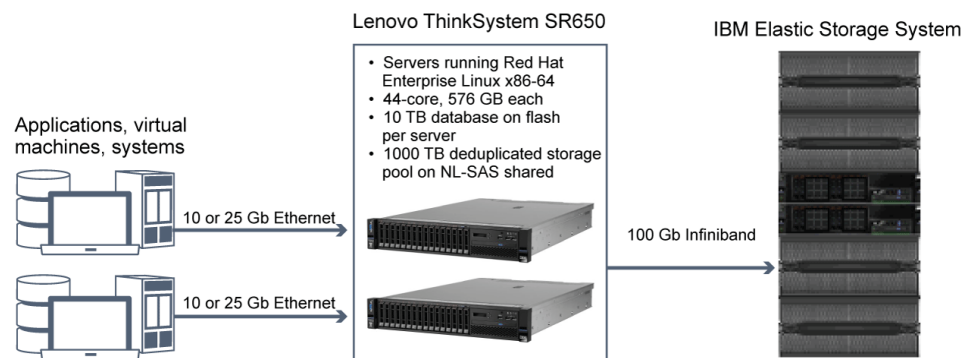


Figure 5. Logical layout for a large IBM Elastic Storage System system

For more information about IBM Elastic Storage System, see the [online product documentation](#).

Chapter 4. Setting up the system

You must set up hardware and preconfigure the system before you run the IBM Spectrum Protect Blueprint configuration script.

About this task

Some steps are unique based on the type of storage that you are configuring for your system. Steps are marked for IBM Flash System storage or IBM Elastic Storage System systems as applicable.

Procedure

1. Configure your storage hardware according to the blueprint specifications and manufacturer instructions. Follow the instructions in [“Step 1: Set up and configure hardware”](#) on page 23.
2. Install the Linux operating system on the server. Follow the instructions in [“Step 2: Install the operating system”](#) on page 25.
3. **IBM Flash System storage:** Configure multipath I/O for disk storage devices. Follow the instructions in [“Step 3: IBM FlashSystem Storage: Configure multipath I/O”](#) on page 28.
4. **IBM FlashSystem Storage:** Create file systems for IBM Spectrum Protect. Follow the instructions in [“Step 4: IBM FlashSystem Storage: Configure file systems for IBM Spectrum Protect”](#) on page 29.
5. **IBM Elastic Storage System:** Configure the IBM Elastic Storage System. Follow the instructions in [“Step 5: IBM Elastic Storage System: Configuring the system”](#) on page 31.
6. Test system performance with the IBM Spectrum Protect workload simulation tool, `tsmdiskperf.pl`. Follow the instructions in [“Step 6: Test system performance”](#) on page 34.
7. Install the IBM Spectrum Protect backup-archive client. Follow the instructions in [“Step 7: Install the IBM Spectrum Protect backup-archive client”](#) on page 37.
8. Install the IBM Spectrum Protect license and server. Follow the instructions in [“Step 8: Install the IBM Spectrum Protect server”](#) on page 38.

Step 1: Set up and configure hardware

Set up and configure storage hardware by using the blueprint specifications and hardware documentation.

About this task

IBM Elastic Storage System: If you are configuring a medium or large system that uses IBM Elastic Storage System, check for system BIOS updates from the server vendor. Apply any suggested changes and then go to [“Step 2: Install the operating system”](#) on page 25.

Procedure

1. Connect your hardware according to manufacturer instructions. For optimal system performance, use at least 8 Gb SAN fabric for connections. If you are using a SAN switch, ensure that it is capable of 8, 16, or 32 Gb connection speeds.
 - For server SAN cabling, use both Fibre Channel connection ports in the dual-port adapters for optimal throughput. Use all four ports in the two dual-port adapters on large systems. All configurations should support a Fibre Channel connection directly to storage or to a SAN switch.
 - For storage subsystem SAN cabling, connect at least two cables to each storage host controller. For a large blueprint, at least four host ports should be cabled on the storage system.
2. Check for system BIOS updates from the server vendor and apply any suggested changes.
3. Configure the disk system.

To configure a IBM FlashSystem disk system, complete the following steps.

Tips:

- For information about using the command line to complete Steps c - e, see [Appendix B, “Configuring the disk system by using commands,”](#) on page 65.
- Small, medium, and large systems were tested by using IBM FlashSystem software level 8.4.0.5.
 - a) Configure licensed functions by following the instructions in your hardware documentation.
 - b) Set up disks in enclosures according to the manufacturer instructions for the size of system that you are building.
 - c) Create RAIDs and LUNs, or volumes. For information about storage configuration layout, see the storage blueprints:
 - [“Small FlashSystem configuration”](#) on page 17
 - [“Medium FlashSystem configuration”](#) on page 18
 - [“Large FlashSystem configuration”](#) on page 20
 - d) Define the IBM Spectrum Protect server as a host to the disk system.
 - e) Assign or map all of the volumes that were created in Step [2c](#) to the new host.

To obtain the Fibre Channel worldwide port name (WWPN) to use for the IBM FlashSystem host mapping, issue the following command:

```
cat /sys/class/fc_host/host*/port_name
```

You should see output similar to the following example:

```
0x10000090fa49009e
0x10000090fa49009f
0x10000090fa3d8f12
0x10000090fa3d8f13
```

If your host is unable to see any devices from the storage system it may be necessary to disable virtualization on one more of the host ports on the IBM FlashSystem.

4. If you attach IBM FlashSystem and IBM Spectrum Protect servers to a SAN fabric, create zones to ensure that specific Fibre Channel ports on the IBM Spectrum Protect server can communicate with specific IBM FlashSystem host ports. During testing, the following guidelines were followed:
 - a. A separate zone was created for each Fibre Channel port on the IBM Spectrum Protect server so that each zone contained no more than one server port.
 - b. Each zone contained one IBM FlashSystem host port from each node canister.

Before you create zones, review the following examples for medium and large systems. The examples are appropriate for a single fabric environment in which the host and disk subsystems are attached to a single switch.

Medium system

- a. On the IBM Spectrum Protect server, both Fibre Channel ports on the dual port Fibre Channel adapter are cabled and are referred to as ha1p1 and ha1p2.
- b. Two of the host ports on the IBM FlashSystem server are cabled (one from each node canister) and are referred to as n1p1 and n2p1.
- c. Two zones are created with the following members:

```
zone1: ha1p1, n1p1, n2p1
zone2: ha1p2, n1p1, n2p1
```

Large system

- a. On the IBM Spectrum Protect server, all four Fibre Channel ports across the two dual port adapters are cabled. The ports are referred to as ha1p1, ha1p2, ha2p1, and ha2p2.
- b. Four of the host ports on the IBM FlashSystem server are cabled (two from each node canister) and are referred to as n1p1, n1p2, n2p1, and n2p2.

c. Four zones are created with the following members:

```
zone1: ha1p1, n1p1, n2p1
zone2: ha1p2, n1p2, n2p2
zone3: ha2p1, n1p1, n2p1
zone4: ha2p2, n1p2, n2p2
```

For additional guidelines about achieving optimal performance and redundancy, see the [SAN configuration and zoning rules summary](#) in [IBM Documentation](#).

Step 2: Install the operating system

Complete the following steps to install Linux x86_64 on the server system.

Before you begin

The operating system is installed on internal hard disk drives. Configure the drives by using a hardware RAID 1 schema. For example, if you are configuring a large system, configure the three 300 GB internal drives by assigning two drives to a RAID 1 pair and the third drive as a spare. If a spare is not needed to meet business requirements, you can configure the system with only two drives.

Procedure

1. Install Red Hat Enterprise Linux Version 7.8 or later or Version 8.5 or later, according to the manufacturer instructions.

Important: Alternatively, you can also choose to install the following operating systems on the server system:

- SUSE Linux Enterprise Server 15 or later version
- Ubuntu 18.04 LTS or later version

Obtain a bootable DVD or .ISO image that contains Red Hat Enterprise Linux at a supported version and start your system from this media. See the following guidance for installation options. If an item is not mentioned in the following list, leave the default selection.

- a) After you start the operating system installation media, choose **Install or upgrade an existing system** from the menu.
- b) On the Welcome screen, select **Test this media & install Red Hat Enterprise Linux 8.x**.
- c) Select your language and keyboard preferences.
- d) Select your location to set the correct timezone.
- e) Select **Software Selection** and then on the next screen, select **Server with GUI**.
- f) From the installation summary page, click **Installation Destination** and verify the following items:
 - The local 300 GB disk is selected as the installation target.
 - Under Other Storage Options, **Automatically configure partitioning** is selected.Click **Done**.
- g) Click **Begin Installation**.

After the installation starts, set the root password for your root user account.

After the installation is completed, restart the system and log in as the root user. Issue the **df** command to verify your basic partitioning.

For example, on a test system, the initial partitioning produced the following result:

```
[root@tvapp02]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root  50G  3.0G  48G   6% /
devtmpfs        32G   0    32G   0% /dev
tmpfs           32G  92K   32G   1% /dev/shm
tmpfs           32G  8.8M  32G   1% /run
```

```
tmpfs          32G     0    32G    0% /sys/fs/cgroup
/dev/mapper/rhel-home 220G   37M  220G    1% /home
/dev/sda1      497M   124M  373M   25% /boot
```

2. Configure your TCP/IP settings according to the operating system installation instructions.

For optimal throughput and reliability, consider bonding multiple network ports together. Bond two ports for a medium system and four ports for a large system. This can be accomplished by creating a Link Aggregation Control Protocol (LACP) network connection, which aggregates several subordinate ports into a single logical connection. The preferred method is to use a bond mode of 802.3ad, **miimon** setting of 100, and a **xmit_hash_policy** setting of layer3+4.

Restriction: To use an LACP network connection, you must have a network switch that supports LACP.

For additional instructions about configuring bonded network connections with Red Hat Enterprise Linux Version 7, see [Create a Channel Bonding Interface](#).

3. Open the `/etc/hosts` file and complete the following actions:

- Update the file to include the IP address and host name for the server. For example:

```
192.0.2.7 server.yourdomain.com server
```

- Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

```
127.0.0.1 localhost
```

4. Install components that are required for the server installation. Complete the following steps to create a Yellowdog Updater Modified (YUM) repository and install the prerequisite packages.

- a) Mount your Red Hat Enterprise Linux installation DVD to a system directory. For example, to mount it to the `/mnt` directory, issue the following command:

```
mount -t iso9660 -o ro /dev/cdrom /mnt
```

- b) Verify that the DVD mounted by issuing the **mount** command.

You should see output similar to the following example:

```
/dev/sr0 on /mnt type iso9660
```

- c) Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

For RHEL 8:

```
cd /etc/yum.repos.d
```

If the `repos.d` directory does not exist, create it.

- d) List directory contents:

```
ls rhel-source.repo
```

- e) Rename the original repo file by issuing the **mv** command.

For example:

```
mv rhel-source.repo rhel-source.repo.orig
```

- f) Create a new repo file by using a text editor.

For example, to use the `vi` editor, issue the following command:

```
vi rhel78_dvd.repo
```

- g) Add the following lines to the new repo file. The **baseurl** parameter specifies your directory mount point:

```
[rhel78_dvd]
name=DVD Redhat Enterprise Linux 7.8
baseurl=file:///mnt
enabled=1
gpgcheck=0
```

For RHEL 8:

```
[InstallMedia-BaseOS]
name=Red Hat Enterprise Linux 8.2.0
mediaid=None
metadata_expire=-1
gpgcheck=0
cost=500
enabled=1
baseurl=file:///mnt/BaseOS/

[InstallMedia-AppStream]
name=Red Hat Enterprise Linux 8.2.0
mediaid=None
metadata_expire=-1
gpgcheck=0
cost=500
enabled=1
baseurl=file:///mnt/AppStream/
```

- h) Install additional prerequisite software packages, by issuing the **yum** command.
For example:

```
yum install ksh.x86_64
yum install sysstat
For RHEL 8:
yum install libnsl
```

5. When the software installation is complete, you can restore the original YUM repository values by completing the following steps:

- a) Unmount the Red Hat Enterprise Linux installation DVD by issuing the following command:

```
umount /mnt
```

- b) Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

- c) Rename the repo file that you created:

```
mv rhel78_dvd.repo rhel78_dvd.repo.orig
```

- d) Rename the original file to the original name:

```
mv rhel-source.repo.orig rhel-source.repo
```

6. Open firewall ports to communicate with the server. Complete the following steps:

- a) Determine the zone that is used by the network interface. The zone is public, by default.

Issue the following command:

```
# firewall-cmd --get-active-zones
public
  interfaces: ens4f0
```

- b) To use the default port address for communications with the server, open TCP/IP port 1500 in the Linux firewall.

Issue the following command:

```
firewall-cmd --zone=public --add-port=1500/tcp --permanent
```

If you want to use a value other than the default, you can specify a number in the range 1024 - 32767. If you open a port other than the default, you will need to specify that port when you run the configuration script.

- c) If you plan to use this system as a hub, open port 11090, which is the default port for secure (https) communications.

Issue the following command:

```
firewall-cmd --zone=public --add-port=11090/tcp --permanent
```

- d) Reload the firewall definitions for the changes to take effect.

Issue the following command:

```
firewall-cmd --reload
```

7. On certain systems, particularly those with more than two CPU sockets, long pauses impacting performance have been observed when the NUMA service attempts to optimize memory. To avoid this, disable the NUMA service.

```
systemctl stop numad  
systemctl disable numad
```

Step 3: IBM FlashSystem Storage: Configure multipath I/O

Complete the steps to enable and configure multipathing for disk storage. For detailed instructions, see the documentation that is provided with your hardware.

Procedure

1. Edit the `/etc/multipath.conf` file to enable multipathing for Linux hosts.

If the `multipath.conf` file does not exist, you can create it by issuing the following command:

```
mpathconf --enable
```

The following parameters were set in `multipath.conf` for testing on an IBM FlashSystem storage system:

```
defaults {  
    user_friendly_names no  
}  
  
devices {  
    device {  
        vendor "IBM "  
        product "2145"  
        path_grouping_policy group_by_prio  
        user_friendly_names no  
        path_selector "round-robin 0"  
        prio "alua"  
        path_checker "tur"  
        failback "immediate"  
        no_path_retry 5  
        rr_weight uniform  
        rr_min_io_rq "1"  
        dev_loss_tmo 120  
    }  
}
```

2. Set the multipath option to start when the system is started.

Issue the following commands:

```
systemctl enable multipathd.service  
systemctl start multipathd.service
```

3. Increase the SCSI timeout for better handling of path failures. For a persistent change, edit the file `/etc/sysconfig/grub` and add the following to the `GRUB_CMDLINE_LINUX` line:

```
scsi_mod.inq_timeout=70
```

Run the following command to rewrite the boot record:

```
grub2-mkconfig -o /etc/grub2.cfg
```

Also, run the following command for an immediate change in addition to the grub update:

```
echo 70 > /sys/module/scsi_mod/parameters/inq_timeout
```

4. To verify that disks are visible to the operating system and are managed by multipath, issue the following command:

```
multipath -l
```

5. Ensure that each device is listed and that it has as many paths as you expect. You can use size and device ID information to identify which disks are listed.

For example, the following output shows that a 2 TB disk has two path groups and four active paths. The 2 TB size confirms that the disk corresponds to a pool file system. Use part of the long device ID number (12, in this example) to search for the volume on the disk-system management interface.

```
[root@tapsrv01 code]# multipath -l
36005076802810c50980000000000012 dm-43 IBM,2145
size=2.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=0 status=active
| |- 2:0:1:18 sdcw 70:64 active undef running
| `-- 4:0:0:18 sdgb 131:112 active undef running
`+- policy='round-robin 0' prio=0 status=enabled
| - 1:0:1:18 sdat 66:208 active undef running
| `-- 3:0:0:18 sddy 128:0 active undef running
```

- a) If needed, correct disk LUN host assignments and force a bus rescan.

For example:

```
echo "- - -" > /sys/class/scsi_host/host0/scan
echo "- - -" > /sys/class/scsi_host/host1/scan
echo "- - -" > /sys/class/scsi_host/host2/scan
```

You can also restart the system to rescan disk LUN host assignments.

- b) Confirm that disks are now available for multipath I/O by reissuing the **multipath -l** command.
6. Use the multipath output to identify and list device IDs for each disk device.

For example, the device ID for your 2 TB disk is 36005076802810c50980000000000012.

Save the list of device IDs to use in the next step.

Step 4: IBM FlashSystem Storage: Configure file systems for IBM Spectrum Protect

You can use the storage preparation script to automate file system configuration or you can complete the process by using manual steps.

About this task

You must format file systems for each of the disk LUNs that the IBM Spectrum Protect server will use. Format database-related file systems and storage pool file systems with xfs.

Note: The /home file system or other file system you have selected for the Db2 instance directory must not be mounted with the nosuid option.

Complete the steps in one of the following sections.

Configure a file system by using the script

To use the configuration script, `storage_prep_lnx.pl`, extract the Blueprint configuration script package and then run the script.

Procedure

1. Open a terminal window and change to the directory where you downloaded the `tsmconfig_v51.tar.gz` file.
2. Extract the file by issuing the following commands:

```
gzip -d tsmconfig_v51.tar.gz
tar -xvf tsmconfig_v51.tar
```

The process creates a directory that is called `tsmconfig`. This directory contains the storage preparation script, the workload simulation tool, and the Blueprint configuration script.

3. Change to the `tsmconfig` directory by issuing the following command:

```
cd tsmconfig
```

4. Run the Perl script and specify the size of system that you are configuring. For example, for a medium system, issue the following command:

```
perl storage_prep_lnx.pl medium
```

If you did not map the disks to the host according to the specifications in “[Step 3: IBM FlashSystem Storage: Configure multipath I/O](#)” on page 28, the script requires customization.

5. List all file systems by issuing the `df` command.

Verify that file systems are mounted at the correct LUN and mount point. Also, verify the available space. The amount of used space should be approximately 1%.

For example:

```
[root@tapsrv04 ~]# df -h /tsminst1/*
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/3600507630081010578000000000003 134G  188M 132G   1%  /tsminst1/
TSMalog
```

Configure a file system by using the manual procedure

You can configure a file system manually by using commands.

Procedure

1. Using the list of device IDs that you generated in “[Step 3: IBM FlashSystem Storage: Configure multipath I/O](#)” on page 28, issue the `mkfs` command to create and format a file system for each storage LUN device. Specify the device ID in the command. For IBM Spectrum Protect V8, format file systems with a command that is similar to the following example:

```
mkfs -t xfs -K /dev/mapper/3600507630081010578000000000002c3
```

Repeat the `mkfs` command for each volume.

2. Create mount point directories for IBM Spectrum Protect file systems.

Issue the `mkdir` command for each directory that you must create. Use the directory values that you recorded in the “[Planning worksheets](#)” on page 10. For example, to create the server instance directory by using the default value, issue the following command:

```
mkdir /tsminst1
```

Repeat the **mkdir** command for each file system.

If you are not using the default paths for your directories, you must manually list directory paths during configuration of the IBM Spectrum Protect server.

3. Add an entry in the `/etc/fstab` file for each file system so that file systems are mounted automatically when the server is started. The entry varies the file system type depending on which file system types were formatted in the previous step.

For example, add the following entry for a XFS file system where the device name is adapted for the actual device name on your system:

```
/dev/mapper/360050763008102618000000000000172 /tsminst1/TSMfile00 xfs
defaults,inode64 0 0
```

4. Mount the file systems that you added to the `/etc/fstab` file by issuing the **mount -a** command.
5. List all file systems by issuing the **df** command.

Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space. The amount of used space should be approximately 1%.

For example:

```
[root@tapshr04 ~]# df -h /tsminst1/*
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/36005076300810105780000000000003 134G  188M 132G  1%  /tsminst1/
TSMalog
```

Step 5: IBM Elastic Storage System: Configuring the system

To configure IBM Elastic Storage System for a large system, review the following information and work with your IBM representative or IBM Business Partner to complete the setup.

Before you begin

This procedure requires IBM Spectrum Scale software:

- Ensure that you can access the installation package for IBM Spectrum Scale Version 5.1.4 or later at the [IBM Passport Advantage®](#) website or at [Fix Central](#).

About this task

Tip: For better integration of your storage environment, install the same level of IBM Spectrum Scale on both the IBM Spectrum Protect server and the IBM Elastic Storage System.

Procedure

1. On the IBM Spectrum Protect system, configure TCP/IP settings according to the manufacturer instructions.
Use different network adapters for communication between server and clients, and the server and the IBM Elastic Storage System.
2. On the IBM Spectrum Protect system, install IBM Spectrum Scale:
 - a) Download the IBM Spectrum Scale base software package at [Passport Advantage](#).
 - b) Download the latest IBM Spectrum Scale fix pack at [Fix Central](#).
 - c) Install the IBM Spectrum Scale base software.
Follow the instructions in [Installing IBM Spectrum Scale on Linux nodes and deploying protocols](#).
 - d) Install the IBM Spectrum Scale fix pack.
3. Ensure that the kernel is portable by issuing the following command:

```
/usr/lpp/mmfs/bin/mmbuildgpl
```

4. Configure a Secure Shell (SSH) automatic login procedure without a password between the IBM Spectrum Protect server and the IBM Elastic Storage System management node and storage nodes. Take one of the following actions:

- If the `/root/.ssh/id_rsa.pub` file is not available on the IBM Spectrum Protect server, generate an `id_rsa.pub` file. The file contains a public key. Issue the following commands from an IBM Elastic Storage System storage node that is part of the cluster:

```
ssh-keygen -t rsa
cd /root/.ssh
chmod 640 /root/.ssh/authorized_keys
```

- If the `/root/.ssh/id_rsa.pub` file is available on the IBM Spectrum Protect server, complete the following steps:
 - a. Append the contents of the `id_rsa.pub` file to the end of the `authorized_keys` file on each of the systems in the IBM Spectrum Scale cluster.
 - b. Append the contents of the `id_rsa.pub` file from each of the other systems in the cluster to the `authorized_keys` file on the IBM Spectrum Protect server.
5. Verify that the login procedure is configured. Log in to the other computers in the cluster from the IBM Spectrum Protect server by running the **ssh** command without using a password.
 6. If the operating system on the IBM Spectrum Protect server is running a firewall, open several ports for incoming network connections from other systems in the IBM Spectrum Scale cluster. For instructions, see [Securing the IBM Spectrum Scale system by using a firewall](#).
 7. Update the `/etc/hosts` file on the IBM Spectrum Scale nodes with information about the IBM Spectrum Protect server.
 8. Add the IBM Spectrum Protect system as an IBM Spectrum Scale node in the cluster by running the **mmaddnode** command. Issue the command from an IBM Elastic Storage System node that is part of the cluster.
For example, if the IBM Spectrum Protect IP address is `192.0.2.7`, you would issue the following command:

```
mmaddnode -N 192.0.2.7
```

9. Assign an IBM Spectrum Scale license to the IBM Spectrum Protect server. From an IBM Elastic Storage System node that is part of the cluster, issue the following command:

```
mmchlicense server -N server_ip_address
```

where `server_ip_address` specifies the IP address of the IBM Spectrum Protect server.

10. To optimize the IBM Spectrum Protect server workload, tune IBM Spectrum Scale client-side parameters by using the **mmchconfig** command.

Issue the following command from an IBM Elastic Storage System node that is part of the cluster:

```
mmchconfig disableDIO=yes,aioSyncDelay=10,pagepool=24G,prefetchAggressivenessRead=0 -N server_ip_address
```

where `server_ip_address` specifies the IP address of the IBM Spectrum Protect server.

If IBM Spectrum Scale replication will be used, the following settings are also required on the IBM Spectrum Protect server to avoid inaccurate capacity reporting.

```
mmchconfig ignoreReplicaSpaceOnStat=yes -i
mmchconfig ignoreReplicationForQuota=yes -i
mmchconfig ignoreReplicationOnStatfs=yes -i
```

11. Create the IBM Spectrum Scale file system on the IBM Elastic Storage System system:
 - a) Verify that the expected factory configuration of a left and right recovery group is in place by using the **mmlsrucoverygroup** command:

- i) Review the command output to verify that two recovery groups exist, and each group has three predefined declustered arrays.
 - ii) Record the recovery group names, which are required in step “11.b” on page 33.
- b) Create a stanza file that defines parameters for each virtual disk:
- i) Specify VDisks in the DA1 declustered array from both recovery groups.
 - ii) Use an 8+2p RAID code for the storage pool data and the 3WayReplication RAID code for the IBM Spectrum Scale file system metadata.

For example, create a file that is named /tmp/ess_vdisk that contains the following information:

```
# cat /tmp/ess_vdisk
%vdisk: vdiskName=GL2_A_L_meta_256k_1 rg=GL2_A_L da=DA1 blocksize=256k
        size=500g raidCode=3WayReplication diskUsage=metadataOnly pool=system
%vdisk: vdiskName=GL2_A_R_meta_256k_1 rg=GL2_A_R da=DA1 blocksize=256k
        size=500g raidCode=3WayReplication diskUsage=metadataOnly pool=system
%vdisk: vdiskName=GL2_A_L_data_8m_1 rg=GL2_A_L da=DA1 blocksize=8m
        raidCode=8+2p diskUsage=dataOnly pool=data
%vdisk: vdiskName=GL2_A_R_data_8m_1 rg=GL2_A_R da=DA1 blocksize=8m
        raidCode=8+2p diskUsage=dataOnly pool=data
```

Because a size is not specified for the two storage pool VDisks, they use all of the remaining space on the declustered arrays.

Tip: For larger file systems, you might have to specify more than two VDisks to meet business requirements. Create VDisks in multiples of 50 TB. Specify the size of the VDisk by using the **SIZE** parameter. For example, to create a 400[®] TB file system, create eight 50 TB VDisks. Stanza entries are similar to the following example:

```
%vdisk: vdiskName=GL2_A_L_data_8m_1
        rg=GL2_A_L da=DA1 blocksize=8m size=50t raidCode=8+2p
        diskUsage=dataOnly pool=data
```

- c) Create disks by running the **mmcrvdisk** and **mmcrnsd** commands and by using the stanza file that you created in step “11.b” on page 33. The **mmcrvdisk** command creates virtual disks, and the **mmcrnsd** command creates IBM Spectrum Scale disks by using the virtual disks.

For example, if the VDisk stanza is called /tmp/ess_vdisk, you would issue the following commands:

```
mmcrvdisk -F /tmp/ess_vdisk
mmcrnsd -F /tmp/ess_vdisk
```

- d) Create a single IBM Spectrum Scale file system by using the **mmcrfs** command and specifying the stanza file. Use the 8 MB block size for data and 256 KB for metadata.
For example:

```
mmcrfs esstsm1 -F /tmp/ess_vdisk -D nfs4 -B 8m --metadata-block-size 256k
        -A yes -L 128M -k nfs4 -m 1 -M 2 -Q no -r 1 -R 2 -S relatime
        -T /esstsm1 -z no
```

- e) Mount the IBM Spectrum Scale file system on the IBM Spectrum Protect system. On the IBM Spectrum Protect system, issue **mmmout** command.
For example:

```
mmmout /esstsm1
```

- f) Verify the amount of free space in the IBM Spectrum Scale file system.
The command and output are similar to the following example:

```
[root@tapsiv03 ~]# df -h /esstsm1
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/esstsm1	401T	1.7G	401T	1%	/esstsm1

- g) Set IBM Spectrum Scale to automatically start when the system starts by using the **chkconfig** command.

For example:

```
chkconfig gpfs on
```

- h) Verify that the VDisks and file system were created correctly by using the **mmlsvdisk** and **mmlsfs** commands.

For example:

```
mmlsvdisk
mmlsfs /dev/esstsm1
```

12. Repeat the above steps to create a file system on the IBM Elastic Storage System 3200 flash storage to be used for the IBM Spectrum Protect database.

What to do next

If you upgrade the Linux operating system to newer kernel levels or you upgrade IBM Spectrum Scale, you must rebuild the portability layer. Follow the instructions in step “3” on page 31.

For more information about completing the steps in the procedure, see the online product documentation:

[Instructions for configuring IBM Elastic Storage System](#)

[Instructions for installing IBM Spectrum Scale](#)

[IBM Spectrum Scale command reference information](#)

Step 6: Test system performance

Before you install the IBM Spectrum Protect server and client, use the workload simulation tool, `tsmdiskperf.pl`, to identify performance issues with your hardware setup and configuration.

About this task

The IBM Spectrum Protect workload simulation tool can test the performance of the IBM Spectrum Protect server database and storage pool disks. The tool, which is a Perl script, uses the **ldeedee** program, which is similar to the Linux operating system **dd** command, to run a non-destructive workload on the system. Use the **iostat** command to monitor the workload for IBM FlashSystem systems.

IBM Elastic Storage System: The tool can report performance statistics only for local devices that are monitored by the **iostat** or **mmpmon** commands. The tool drives loads against other network-attached devices, but does not collect and report on performance statistics. When the tool is run against a file system on an IBM Elastic Storage System, the tool automatically runs the **mmpmon** command.

Sample data from the **iostat** command is extracted for the specific disks that were involved in the test. Then, peak and average measurements for input/output operations per second (IOPS) and throughput are calculated. The script uses the **ldeedee** command across multiple threads to drive the I/O by using direct I/O.

Tips:

- The **iostat** tool monitors and reports on all I/O for the related disks, even activity that is being driven by applications other than the workload tool. For this reason, ensure that other activity is stopped before you run the tool.
- New storage arrays go through an initialization process. Allow this process to end before you measure disk performance. On IBM FlashSystem disk systems, you can monitor the initialization progress in the **Running Tasks** view.

The workload simulation tool can run the following types of workloads:

Storage pool workload

The storage pool workload simulates IBM Spectrum Protect server-side data deduplication, in which large, 256 KB block-size sequential read and write operations are overlapped. The write process simulates incoming backups while the read operation simulates identification of duplicate data. The tool creates a read and write thread for every file system that is included in the test, allowing multiple sessions and processes to be striped across more than one file system.

You can also simulate a storage pool workload that conducts only read I/O or only write I/O operations:

- Simulate restore operations by specifying the **mode=readonly** option.
- Simulate backup operations by specifying the **mode=writeonly** option.

Database workload

The database workload simulates IBM Spectrum Protect database disk access in which small, 8 KB read and write operations are performed randomly across the disk. For this workload, 10 GB files are pre-created on each of the specified file systems and then read and write operations are run to random ranges within these files. Multiple threads are issued against each file system, sending I/O requests simultaneously.

For the database workload, configurations typically have one file system for each pool on the storage array. Include all database file systems when you are testing the database workload.

To use the tool effectively, experiment with test runs by including different quantities of file systems in the simulation until the performance of the system diminishes.

Depending on disk speed and the number of file systems that you are testing, the time that is required to run the script can be 3 - 10 minutes.

Procedure

To use the workload simulation tool, complete the following steps:

1. Plan to test either the storage pool file systems or the database file systems.
2. Collect a list of the file systems that are associated with your chosen type of storage. Break the file systems into groups according to which pool they belong to on the disk system.

Grouping is used to ensure that physical disks from all volumes on all arrays for the storage type are engaged in the test. To review groupings for file systems, see the volume configuration tables in Chapter 3, “Storage configuration blueprints,” on page 15.

IBM Elastic Storage System: Because only a single IBM Spectrum Scale file system is defined for storage, you must create temporary directories to use when you run the workload simulation tool and specify the **-fslist** option. For example, issue the **mkdir** command to create temporary directories:

```
mkdir /esstsm1/perftest/1
mkdir /esstsm1/perftest/2
< ... >
mkdir /esstsm1/perftest/14
```

3. To run the tool, change to the **tsmconfig** directory by issuing the following command:

```
cd tsmconfig
```

If you did not extract the Blueprint configuration script compressed file to prepare file systems for IBM Spectrum Protect, follow the instructions in “Configure a file system by using the script” on page 30.

4. Run an initial test of the workload that includes one file system of the storage type from each pool on the storage array.
For example, to simulate the IBM Spectrum Protect storage pool workload on a medium-scale system, issue the following command:

```
perl tsmdiskperf.pl workload=stgpool
fslist=/tsminst1/TSMfile00,/tsminst1/TSMfile01,/tsminst1/TSMfile02,/tsminst1/
```

```
TSMfile03,/tsminst1/TSMfile04,/tsminst1/TSMfile05,/tsminst1/TSMfile06,/tsminst1/
TSMfile07
```

For example, to simulate backup operations (by using only write I/O) for an IBM Spectrum Protect storage pool workload on a medium-scale system, issue the following command:

```
perl tsmdiskperf.pl workload=stgpool
fslist=/tsminst1/TSMfile00,/tsminst1/TSMfile01,/tsminst1/TSMfile02,/tsminst1/
TSMfile03,/tsminst1/TSMfile04,/tsminst1/TSMfile05,/tsminst1/TSMfile06,/tsminst1/
TSMfile07 mode=writeonly
```

To simulate the database workload on a small-scale system and include all four of the database file systems, issue the following command:

```
perl tsmdiskperf.pl workload=db fslist=/tsminst1/TSMdbspace00,/tsminst1/
TSMdbspace01,/tsminst1/TSMdbspace02,/tsminst1/TSMdbspace03
```

Record the reported results for each test run.

5. If you have implemented a storage configuration with multiple arrays that are not combined into single storage pool, rerun the previous test, but modify it to include one additional file system from each pool.
For example, if you have two pools on the array that is dedicated to the storage pool, your test sequence will include a count of file systems at 2, 4, 6, 8, 10, and so on.
6. Continue repeating these tests while the reported performance measurements improve. When performance diminishes, capture the results of the last test that indicated improvement. Use these results as the measurements for comparison.

Results

The performance results that are provided when you run the workload simulation tool might not represent the maximum capabilities of the disk subsystem that is being tested. The intent is to provide measurements that can be compared against the lab results that are reported for medium and large systems.

The workload simulation tool is not intended to be a replacement for disk performance analysis tools. Instead, you can use it to spot configuration problems that affect performance before you run IBM Spectrum Protect workloads in a production environment. Problems will be evident if the measurements from test runs are significantly lower than what is reported for test lab systems. If you are using hardware other than the Storwize® components that are included in this document, use your test results as a rough estimate of how other disk types compare with the tested configurations.

Example

This example shows the output from a storage pool workload test on a small system. Eight file systems are included. The following command is issued:

```
perl tsmdiskperf.pl workload=stgpool fslist=/tsminst1/TSMfile00,/tsminst1/TSMfile01,/
tsminst1/TSMfile02,/tsminst1/TSMfile03,/tsminst1/TSMfile04,/tsminst1/TSMfile05,/
tsminst1/TSMfile06,/tsminst1/TSMfile07
```

The output shows the following results:

```
=====
: IBM Spectrum Protect disk performance test      (Program version 5.1)
:
: Workload type:                                stgpool
: Number of filesystems:                        8
: Mode:                                         readwrite
: Files to write per fs:                        5
: File size:                                    2 GB
:
=====
:
```

```

: Beginning I/O test.
: The test can take upwards of ten minutes, please be patient ...
: Starting write thread ID: 1 on filesystem /tsminst1/TSMfile00
: Starting read thread ID: 2 on filesystem /tsminst1/TSMfile00
: Starting write thread ID: 3 on filesystem /tsminst1/TSMfile01
: Starting read thread ID: 4 on filesystem /tsminst1/TSMfile01
: Starting write thread ID: 5 on filesystem /tsminst1/TSMfile02
: Starting read thread ID: 6 on filesystem /tsminst1/TSMfile02
: Starting write thread ID: 7 on filesystem /tsminst1/TSMfile03
: Starting read thread ID: 8 on filesystem /tsminst1/TSMfile03
: Starting write thread ID: 9 on filesystem /tsminst1/TSMfile04
: Starting read thread ID: 10 on filesystem /tsminst1/TSMfile04
: Starting write thread ID: 11 on filesystem /tsminst1/TSMfile05
: Starting read thread ID: 12 on filesystem /tsminst1/TSMfile05
: Starting write thread ID: 13 on filesystem /tsminst1/TSMfile06
: Starting read thread ID: 14 on filesystem /tsminst1/TSMfile06
: Starting write thread ID: 15 on filesystem /tsminst1/TSMfile07
: Starting read thread ID: 16 on filesystem /tsminst1/TSMfile07
: All threads are finished. Stopping iostat process with id 15732
=====
: RESULTS:
: Devices reported on from output:
: dm-25
: dm-28
: dm-7
: dm-6
: dm-4
: dm-8
: dm-12
: dm-15
:
: Average R Throughput (KB/sec):      227438.06
: Average W Throughput (KB/sec):      224826.38
: Avg Combined Throughput (MB/sec):    441.66
: Max Combined Throughput (MB/sec):    596.65
:
: Average IOPS:                        1767.16
: Peak IOPS:                           2387.43 at 08/05/2015 09:38:27
:
: Total elapsed time (seconds):         171
=====

```

What to do next

Compare your performance results against test lab results by reviewing sample outputs for storage pool and database workloads on both medium and large systems:

- For the storage pool workload, the measurement for average combined throughput in MB per second combines the read and write throughput. This is the most useful value when you compare results.
- For the database workload, the peak IOPS measurements add the peak read and write operations per second for a specific time interval. This is the most useful value when you compare results for the database workload.

To review the sample outputs, see [Appendix A, “Performance results,”](#) on page 57.

Step 7: Install the IBM Spectrum Protect backup-archive client

Install the IBM Spectrum Protect for Linux backup-archive client so that the administrative command-line client is available.

About this task

Install the backup-archive client and API on the server system.

Procedure

1. Change to the directory where you downloaded the client package files.

2. Install the software by issuing the appropriate commands.

For example:

```
yum install ./gskcrypt64-8.0.55.24.linux.x86_64.rpm  
./gskssl64-8.0.55.24.linux.x86_64.rpm  
yum install ./TIVsm-API64.x86_64.rpm ./TIVsm-BA.x86_64.rpm
```

For detailed installation instructions, see [Installing the backup-archive clients](#) in IBM Knowledge Center.

Tip: If available, you can display different versions of the same topic by using the versions menu at the top of the page.

Step 8: Install the IBM Spectrum Protect server

Before you can run the Blueprint configuration script, you must install the IBM Spectrum Protect server and license.

About this task

Install IBM Spectrum Protect V8.1.14.100 or a later level. In this way, you can obtain critical fixes. To take advantage of the latest product updates, install the latest product level. At the time of publication, the latest available level was V8.1.16.

Before you install IBM Spectrum Protect, review the list of new features, including any security enhancements, for your selected release. For an overview, see [What's new in V8 releases](#).

For information about security updates, see [What you should know about security before you install or upgrade the server](#).

Obtain the installation package

You can obtain the IBM Spectrum Protect installation package from an IBM download site such as Passport Advantage or IBM Fix Central.

Before you begin

To ensure that the files can be downloaded correctly, set the system user limit for maximum file size to unlimited:

1. To query the maximum file size value, issue the following command:

```
ulimit -Hf
```

2. If the system user limit for maximum file size is not set to unlimited, change it to unlimited by following the instructions in the documentation for your operating system.

Procedure

1. Download the server installation package from [Passport Advantage](#) or [Fix Central](#).
2. For the latest information, updates, and maintenance fixes, go to the [IBM Support Portal](#).
3. Complete the following steps:
 - a) Verify that you have enough space to store the installation files when they are extracted from the product package. See the download document for the space requirements:
 - IBM Spectrum Protect: [technote 4042992](#)
 - IBM Spectrum Protect Extended Edition: [technote 4042992](#)
 - b) Download the package to the directory of your choice. The path must contain no more than 128 characters. Be sure to 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.

- c) Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name
```

where *package_name* is the name of the downloaded package, for example:

```
8.1.16.000-IBM-SPSRV-Linuxx86_64.bin
```

- d) Extract the file from the package by issuing the following command:

```
./package_name
```

where *package_name* is the name of the package.

Install the IBM Spectrum Protect server

Install IBM Spectrum Protect V8.1.12 or later, by using the command line in console mode. To take advantage of product updates and fixes for possible issues, install the latest level. At the time of publication, the latest level of IBM Spectrum Protect was 8.1.16.000.

Before you begin

Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

During the IBM Spectrum Protect installation process, you must temporarily disable the Security-Enhanced Linux (SELinux) security module if it is enabled. If SELinux is enabled, the IBM Spectrum Protect installation process fails. Complete the following steps:

1. Determine whether SELinux is installed and in enforcing mode by taking one of the following actions:
 - Check the `/etc/sysconfig/selinux` file.
 - Run the **sestatus** operating system command.
 - Check the `/var/log/messages` file for SELinux notices.
2. If SELinux is enabled, disable it (or set it to permissive) by completing the following steps:
 - a. Set permissive mode by issuing the **setenforce 0** command as a superuser.
 - b. Modify the `/etc/sysconfig/selinux` file to set `SELINUX=disabled`.

Procedure

To install IBM Spectrum Protect, complete the following steps:

1. Change to the directory where you downloaded the package.
2. Start the installation wizard in console mode by issuing the following command:

```
./install.sh -c
```

Optional: Generate a response file as part of a console mode installation. Complete the console mode installation options, and in the **Summary** window, specify **G** to generate the responses.

Results

If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory, for example:

```
/var/ibm/InstallationManager/logs
```

What to do next

After the installation finishes, re-enable SELinux.

Before you customize IBM Spectrum Protect for your use, go to the [IBM Support Portal for IBM Spectrum Protect](#). Click **Downloads (fixes and PTFs)** and apply any applicable fixes.

Tip: For more information about installation, see [Installing the server components](#) in IBM Knowledge Center.

Chapter 5. Configuring the IBM Spectrum Protect server

Run the Blueprint configuration script, `TSMserverconfig.pl`, to configure the IBM Spectrum Protect server.

Before you begin

You can run the Blueprint configuration script in interactive or non mode. In interactive mode, you provide responses for each step in the script and accept defaults or enter values for the configuration. In noninteractive mode, the script uses a response file that contains answers to the script prompts.

To run the script in noninteractive mode, use one of the response files that are included in the blueprint configuration compressed file. For instructions about how to use a response file, see [Appendix C, “Using a response file with the Blueprint configuration script,” on page 71](#).

About this task

When you start the script and select the size of the server that you want to configure, the script verifies the following hardware and system configuration prerequisites:

- Sufficient memory is available for server operations.
- Processor core count meets blueprint specifications.
- Kernel parameters are set correctly. If the values are not set as specified, they are automatically updated when you run the Blueprint configuration script to configure the server. For more information about kernel parameter settings, see [Table 20 on page 47](#).
- All required file systems are created.
- The minimum number of file system types exist and the minimum level of free space is available in each file system.

If all prerequisites checks are passed, the script begins server configuration. The following tasks are completed to configure the server for optimal performance, based on the scale size that you select:

- A Db2 database instance is created.
- The `dsmserv.opt` options file with optimum values is created.
- The server database is formatted.
- The system configuration is updated to automatically start the server when the system starts.
- Definitions that are required for database backup operations are created.
- A directory-container storage pool with optimal performance settings for data deduplication is defined.

You can use the `-legacy` option with the blueprint configuration script to force the creation of a deduplicated storage pool, which uses a FILE device class.

- Policy domains for each type of client workload are defined.
- Schedules for client backup are created.
- Server maintenance schedules that are sequenced for optimal data deduplication scalability are created.
- The client options file is created.

The blueprint configuration script includes a compression option that enables compression for both the archive log and database backups. You can save significant storage space by using this option, but the amount of time that is needed to complete database backups increases. The preferred method is to enable the option if you are configuring a small blueprint system because limited space is configured for the archive log and database backups.

The default setting for the compression option is disabled.

Tip: Do not confuse the blueprint configuration script compression option with inline compression of data in container storage pools, which is enabled by default with IBM Spectrum Protect V7.1.5 and later.

Complete the following steps as the root user to run the Blueprint configuration script.

Procedure

1. Open a terminal window.
2. If you did not extract the Blueprint configuration script compressed file to prepare file systems for IBM Spectrum Protect, follow the instructions in [“Configure a file system by using the script”](#) on page 30.
3. Change to the `tsmconfig` directory by issuing the following command:

```
cd tsmconfig
```

4. Run the configuration script in one of the following modes:

- To run the configuration script in interactive mode and enter your responses at the script prompts, issue the following command:

```
perl TSMserverconfig.pl
```

If you want to enable compression for the archive log and database backups on a small system, issue the following command:

```
perl TSMserverconfig.pl -compression
```

Depending on how you preconfigured the system, you can accept the default values that are presented by the script. Use the information that you recorded in the [“Planning worksheets”](#) on page 10 as a guide. If you changed any of the default values during the preconfiguration step, you must manually enter your values at the script prompts.

- To run the configuration script in noninteractive mode by using a response file to set configuration values, specify the response file when you run the script. For example:
 - To use the default response file for a medium system, issue the following command:

```
perl TSMserverconfig.pl responsefilemed.txt
```

- To use the default response file for a small system and enable compression for the archive log and database backups, issue the following command:

```
perl TSMserverconfig.pl responsefilesmall.txt -compression
```

- To use the default response file for a system that uses IBM Elastic Storage System, issue the following command:

```
perl TSMserverconfig.pl responsefile_ess.txt -skipmount
```

If you encounter a problem during the configuration and want to pause temporarily, use the quit option. When you run the script again, it resumes at the point that you stopped. You can also open other terminal windows to correct any issues, and then return to and continue the script. When the script finishes successfully, a log file is created in the current directory.

5. Save the log file for future reference.

The log file is named `setupLog_datestamp.log` where *datestamp* is the date on which you ran the configuration script. If you run the script more than once on the same day, a version number is appended to the end of the name for each additional version that is saved.

For example, if you ran the script three times on July 27, 2013, the following logs are created:

- `setupLog_130727.log`
- `setupLog_130727_1.log`

- `setupLog_130727_2.log`

Results

After the script finishes, the server is ready to use. Review [Table 19 on page 43](#) and the setup log file for details about your system configuration.

[Table 20 on page 47](#) provides details about kernel parameter values for the system. Also, consider tuning the `TCPWINDOWSIZE` option to 0 for Linux servers and clients.

Table 19. Summary of configured elements	
Item	Details
Db2 database instance	<ul style="list-style-type: none"> • The Db2 instance is created by using the instance user ID and instance home directory. • Db2 instance variables that are required by the server are set. • The Db2 -locklist parameter remains at the default setting of Automatic (for automatic management), which is preferred for container storage pools. If you are defining a non-container storage pool, you can use the -locklist parameter with the IBM Spectrum Protect blueprint configuration script, TSMserverconfig.pl, to revert to manually setting -locklist values.
Operating system user limits (ulimits) for the instance user	<p>The following values are set:</p> <ul style="list-style-type: none"> • Maximum size of core files created (<code>core</code>): unlimited • Maximum size of a data segment for a process (<code>data</code>): unlimited • Maximum file size allowed (<code>fsize</code>): unlimited • Maximum number of open files that are allowed for a process (<code>nofile</code>): 65536 • Maximum amount of processor time in seconds (<code>cpu</code>): unlimited • Maximum number of user processes (<code>nproc</code>): 16384
IBM Spectrum Protect API	<ul style="list-style-type: none"> • An API <code>dsm.sys</code> file is created in the <code>/opt/tivoli/tsm/server/bin/dbbkapi/</code> directory. The following parameters are set. Some values might vary, depending on selections that were made during the configuration: <ul style="list-style-type: none"> servername <code>TSMDBMGR_tsminst1</code> tcpserveraddr <code>localhost</code> commmethod <code>tcpip</code> tcpserveraddr <code>localhost</code> tcpport <code>1500</code> passworddir <code>/home/tsminst1/tsminst1</code> errorlogname <code>/home/tsminst1/tsminst1/tsmdbmgr.log</code> nodename <code>\$\$_TSMDBMGR_\$\$</code> • The API password is set.
Server settings	<ul style="list-style-type: none"> • The server is configured to start automatically when the system is started. • An initial system level administrator is registered. • The server name and password are set. • The following values are specified for SET commands: <ul style="list-style-type: none"> – SET ACTLOGRETENTION is set to 180. – SET EVENTRETENTION is set to 180. – SET SUMMARYRETENTION is set to 180.

Table 19. Summary of configured elements (continued)

Item	Details
IBM Spectrum Protect server options file	<p>The <code>dsmserv.opt</code> file is set with optimal parameter values for server scale. The following server options are specified:</p> <ul style="list-style-type: none"> • ACTIVELOGSIZE is set according to scale size: <ul style="list-style-type: none"> – Extra Small system: 24576 – Small system: 131072 – Medium system: 131072 – Large system: 524032 • If you enabled compression for the blueprint configuration, ARCHLOGCOMPRESS is set to Yes. • COMMTIMEOUT is set to 3600 seconds. • If you are using the <code>-legacy</code> option for data deduplication, DEDUPDELETIONTHREADS is set according to scale size: <ul style="list-style-type: none"> – Extra Small system: 2 – Small system: 8 – Medium system: 8 – Large system: 12 • DIOENABLED is set to NO for IBM Elastic Storage System configurations when a directory-container storage pool is created. • DIRECTIO is set to NO for IBM Elastic Storage System configurations. For Storwize configurations, the preferred method is to use the default value of YES. • DEDUPREQUIRESBACKUP is set to NO. • DEVCONFIG is specified as <code>devconf.dat</code>, which is where a backup copy of device configuration information will be stored. • EXPINTERVAL is set to 0, so that expiration processing runs according to schedule. • IDLETIMEOUT is set to 60 minutes. • MAXSESSIONS is set according to scale size: <ul style="list-style-type: none"> – Extra Small system: 75 maximum simultaneous client sessions – Small system: 250 maximum simultaneous client sessions – Medium system: 500 maximum simultaneous client sessions – Large system: 1000 maximum simultaneous client sessions <p>The effective value for the SET MAXSCHEDESESSIONS option is 80% of the value that was specified for the MAXSESSIONS option:</p> <ul style="list-style-type: none"> – Extra Small system: 45 sessions (60%) – Small system: 200 sessions – Medium system: 400 sessions – Large system: 800 sessions • NUMOPENVOLSAALLOWED is set to 20 open volumes. • TCPWINDOWSIZE is set to 0 • VOLUMEHISTORY is specified as <code>volhist.dat</code>, which is where the server will store a backup copy of volume history information. In addition to <code>volhist.dat</code>, which will be stored in the server instance directory, a second volume history option is specified to be stored in the first database backup directory for redundancy.

Table 19. Summary of configured elements (continued)

Item	Details
IBM Spectrum Protect server options file: database reorganization options	<p>Server options that are related to database reorganization are specified in the following sections.</p> <p>Servers at V7.1.1 or later:</p> <ul style="list-style-type: none"> • ALLOWREORGINDEX is set to YES. • ALLOWREORGTABLE is set to YES. • DISABLEREORGINDEX is not set. • DISABLEREORGTABLE is set to <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <p>BF_AGGREGATED_BITFILES, BF_BITFILE_EXTENTS, ARCHIVE_OBJECTS, BACKUP_OBJECTS</p> </div> <ul style="list-style-type: none"> • REORGBEGINTIME is set to 12:00. • REORGDURATION is set to 6.
Directory-container storage pool	<p>A directory-container storage pool is created, and all of the storage pool file systems are defined as container directories for this storage pool. The following parameters are set in the DEFINE STGPPOOL command:</p> <ul style="list-style-type: none"> • STGTYPE is set to DIRECTORY. • MAXWRITERS is set to NOLIMIT. <p>For servers at V7.1.5 or later, compression is automatically enabled for the storage pool.</p>
Storage pool if the -legacy option is specified	<ul style="list-style-type: none"> • A FILE device class is created and tuned for configuration size: <ul style="list-style-type: none"> – All storage pool file systems are listed with the DIRECTORY parameter in the DEFINE DEVCLASS command. – The MOUNTLIMIT parameter is set to 4000 for all size systems. – The MAXCAP parameter is set to 50 GB for all size systems. • The storage pool is created with settings that are tuned for configuration size: <ul style="list-style-type: none"> – Data deduplication is enabled. – The value of the IDENTIFYPROCESS parameter is set to 0 so that duplicate identification can be scheduled. – Threshold reclamation is disabled so that it can be scheduled. – The MAXSCRATCH parameter value is tuned based on the amount of storage that is available in the FILE storage pool.

Table 19. Summary of configured elements (continued)

Item	Details
Server schedules	<p>The following server maintenance schedules are defined:</p> <ul style="list-style-type: none"> A replication storage rule is scheduled to run 10 hours after the start of the backup window. The schedule is inactive by default. You must specify the parameter ACTIVE=Yes to enable the processing of the replication storage rule at the scheduled time. <p>Remember: If a replication storage rule is configured with the parameter ACTIONTYPE=NOREPLICATING, then you must define a replication subrule for the parent replication storage rule with the parameter ACTIONTYPE=REPLICATE to replicate data from specific nodes and filesystem.</p> <p>Sessions are based on system size:</p> <ul style="list-style-type: none"> Extra Small system: 8 Small system: 20 Medium system: 40 Large system: 60 <ul style="list-style-type: none"> Database backup is scheduled to run until it is complete. The schedule starts 14 hours after the beginning of the client backup window. <p>A device class that is named DBBACK_FILEDEV is created for the database backup. If the configuration script is started with the compression option, the BACKUP DB command runs with compress=yes.</p> <p>The device class is created to allow a mount limit of 32. The file volume size is set to 50 GB. The device class directories include all of the database backup directories. The number of database backup sessions is based on the system size:</p> <ul style="list-style-type: none"> Extra Small system: 2 Small system: 2 Medium system: 4 Large system: 12 <p>In addition, the SET DBRECOVERY command is issued. It specifies the device class, the number of streams, and the password for database backup operations. After a successful database backup operation, the DELETE VOLHISTORY command is used to delete backups that were created more than 4 days prior.</p> <ul style="list-style-type: none"> Expiration processing is scheduled to run until it is complete. The schedule starts 17 hours after the beginning of the client backup window. The RESOURCE parameter is set according to scale size and type of data deduplication storage pool: <p>Directory-container storage pools:</p> <ul style="list-style-type: none"> Extra Small system: 4 Small system: 10 Medium system: 30 Large system: 40 <p>Non-container storage pools:</p> <ul style="list-style-type: none"> Small system: 6 Medium system: 8 Large system: 10 <p>If you are using the -legacy option for data deduplication, the following schedules are also defined:</p> <ul style="list-style-type: none"> Duplicate identification is set for a duration of 12 hours. The schedule starts at the beginning of the client backup window. The NUMPROCESS parameter is set according to scale size: <ul style="list-style-type: none"> Extra Small system: 1 Small system: 12 Medium system: 16 Large system: 32 Reclamation processing is set for a duration of 8 hours. The reclamation threshold is 25%. The schedule starts 14 hours after the beginning of the client backup window. The RECLAIMPROCESS parameter is set as part of the storage pool definition, according to scale size: <ul style="list-style-type: none"> Extra Small system: 2 Small system: 10 Medium system: 20 Large system: 32

Table 19. Summary of configured elements (continued)	
Item	Details
Policy domains	<p>The following policy domains are created:</p> <ul style="list-style-type: none"> • STANDARD – The default policy domain • <i>server name</i>_DATABASE – Policy domain for database backups • <i>server name</i>_DB2 – Policy domain for Db2 database backups • <i>server name</i>_FILE – Policy domain for file backups that use the backup-archive client • <i>server name</i>_MAIL – Policy domain for mail application backups • <i>server name</i>_ORACLE – Policy domain for Oracle database backups • <i>server name</i>_VIRTUAL – Policy domain for virtual machine backups • <i>server name</i>_HANA – Policy domain for SAP HANA backups • <i>server name</i>_OBJECT - Policy domain for Amazon Simple Storage Service (S3) object data from IBM Spectrum Protect Plus offload operations <p>Policy domains other than the STANDARD policy are named by using a default value with the server name. For example, if your server name is TSMSERVER1, the policy domain for database backups is TSMSERVER1_DATABASE.</p>
Management classes	<p>Management classes are created within the policy domains that are listed in the previous row. Retention periods are defined for 7, 30, 90, and 365 days.</p> <p>The default management class uses the 30-day retention period.</p>
Client schedules	<p>Client schedules are created in each policy domain with the start time that is specified during configuration.</p> <p>The type of backup schedule that is created is based on the type of client:</p> <ul style="list-style-type: none"> • File server schedules are set as incremental forever. • Data protection schedules are set as full daily. <p>Some data protection schedules include command file names that are appropriate for the data protection client.</p> <p>For more information about the schedules that are predefined during configuration, see Appendix D, “Using predefined client schedules,” on page 73.</p>

Table 20. Linux kernel parameter optimum settings		
Parameter	Description	Preferred value
kernel.shmmni, kernel.shmmax, kernel.shmall, kernel.sem, kernel.msgmni, kernel.msgmax, kernel.msgmnb		These values are set automatically by the Db2 software, as described in Modifying kernel parameters (Linux) .
kernel.randomize_va_space	The kernel.randomize_va_space parameter configures the use of memory address space layout randomization (ASLR) for the kernel. If you set the value to 0, you disable ASLR. To learn more about Linux ASLR and Db2, see technote 1365583 .	Set this parameter value to 2, which is the default value for the operating system. Later, if you decide not to use ASLR, you can reset the value to 0.
vm.swappiness	The vm.swappiness parameter defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the Db2 product information in IBM Documentation .	If you installed an IBM Spectrum Protect V8 server, set this parameter to 5. If you installed a V7 server, set this parameter to 0.
vm.overcommit_memory	The vm.overcommit_memory parameter influences how much virtual memory can be allocated, based on kernel parameter settings. For more information about kernel parameters, see the Db2 product information in IBM Documentation .	0

Removing an IBM Spectrum Protect blueprint configuration

If your blueprint configuration fails, you can use a cleanup script to remove the IBM Spectrum Protect server and stored data.

Before you begin



Attention: The automated script `cleanupserversetup.pl` is destructive and will completely remove an IBM Spectrum Protect server and all stored data.

About this task

The script can be used to clean up your system during initial testing and troubleshooting for blueprint configurations. If a configuration attempt fails, running the script removes the server and all associated IBM Spectrum Protect data. The script uses the file, `serversetupstatefileforcleanup.txt`, which is generated when you run the Blueprint configuration script, `TSMserverconfig.pl`.

The cleanup script is available in the `diag` folder of the blueprint `tsmconfig` directory.

Procedure

To clean up your system by using the script, complete the following steps:

1. Edit the `cleanupserversetup.pl` script by commenting out the exit on the first line.

For example:

```
#exit;    # This script is destructive, so by default it exits. Comment-out  
          this line to proceed.
```

2. Copy the `cleanupserversetup.pl` script into the folder where the `TSMserverconfig.pl` script is located.
3. Issue the following command:

```
perl cleanupserversetup.pl
```

Chapter 6. Completing the system configuration

Complete the following tasks after your IBM Spectrum Protect server is configured and running.

About this task

For more information about the configuration tasks, see the documentation for your IBM Spectrum Protect server version in [IBM Knowledge Center](#).

Tip: To display a different version of the same topic in IBM Knowledge Center, you can use the versions menu, if available, at the top of the page.

Changing initial passwords

Beginning with v5.1, the blueprint scripts no longer allow the use of a default password. The user is required to provide a password during the configuration in response to prompts when running interactively, or in the response files when running non-interactively.

About this task

The script sets as specified by the user for the following passwords:

- Initial IBM Spectrum Protect administrator
- IBM Spectrum Protect server
- Db2 instance owner

Procedure

- To update password information for the server and administrator, use server commands. For more information, see the **SET SERVERPASSWORD**, **UPDATE ADMIN**, and **UPDATE SERVER** server commands.
- To update the password for the Db2 instance owner, use the Linux operating system **passwd** command.
- Create a system-level administrator. Then, remove or lock the administrator that is named ADMIN by using the **REMOVE ADMIN** or **LOCK ADMIN** command.
- Change the password that is used to protect the server encryption key for database backup operations. Issue the following command:

```
set dbrecovery dbback_filedev password=<newpassword>
```

where *newpassword* is the password that you set.



Attention: You must remember the password, or you will be unable to restore database backups.

Registering nodes and associating them with predefined client schedules

When you are ready to register nodes to the IBM Spectrum Protect server, use the **REGISTER NODE** command. Then, you can associate nodes with a predefined client schedule.

Before you begin

When you register nodes, the host name of the protected system is typically used for the node name. In the following example, assume that you want to register a node named *newnode1* to the

TSMSERVER1_FILE domain for backup-archive client backups, and associate the node with a predefined client schedule. You can use the administrative command line to issue server commands for the operation.

When you issue the **REGISTER NODE** server command, increase the default value for the maximum number of mount points that a node is allowed to use on the server. Specify a value of 99 for the **MAXNUMMP** parameter instead of using the default.

Complete the following example steps to register *newnode1*, associate it with a schedule, and then verify that the schedule is ready to use for backups.

Procedure

1. Register *newnode1* to the TSMSERVER1_FILE domain. Specify a value for the client node password, for example, *pw4node1*. Set the **MAXNUMMP** parameter to 99:

```
register node newnode1 pw4node1 dom=TSMSERVER1_FILE maxnummp=99
```

2. To use a predefined client schedule, determine which schedule to associate *newnode1* with by querying the list of available schedules. Issue the **QUERY SCHEDULE** command.

The output lists all defined schedules. For example, the following output shows the details for the FILE _INCRFOREVER_10PM schedule:

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
TSMSERVER1_FILE	FILE_INCRFOREVER_10PM	Inc Bk	07/24/2013 22:00:00	60 M	1 D	Any

3. Define an association between *newnode1* and the FILE _INCRFOREVER_10PM schedule. You must specify the domain for the node and schedule.
For example:

```
define association TSMSERVER1_FILE FILE_INCRFOREVER_10PM newnode1
```

4. Verify that *newnode1* is associated with the correct schedule by issuing the **QUERY ASSOCIATION** command.

For example, issue the following command, specifying the schedule domain and the schedule name:

```
query association TSMSERVER1_FILE FILE_INCRFOREVER_10PM
```

The output shows that *newnode1* is associated with the queried domain and schedule name.

```
Policy Domain Name: TSMSERVER1_FILE
Schedule Name: FILE_INCRFOREVER_10PM
Associated Nodes: NEWNODE1
```

5. Display details about the client schedule by issuing the **QUERY EVENT** command. Specify the domain and name of the schedule for which you want to display events.

For example, issue the following command:

```
query event TSMSERVER1_FILE FILE_INCRFOREVER_10PM
```

The output shows that the backup for *newnode1* is scheduled, but has not yet occurred.

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
08/23/2013 22:00:00		FILE_INCRFOREVER_10PM	NEWNODE1	Future

6. After you register a node and assign it to a schedule, configure the client and client schedule on the client system and then start the scheduler daemon on the client system so that the backup operation starts at the scheduled time.

To configure the client schedules that are predefined by the Blueprint configuration script, see [Appendix D, "Using predefined client schedules,"](#) on page 73.

For more information about starting the client scheduler, see the IBM Spectrum Protect [client documentation](#) in IBM Knowledge Center.

Reorganizing database tables and indexes

Schedule database table and index reorganization to ensure that the server is running efficiently.

About this task

If tables or the indexes that are associated with tables are not reorganized, unexpected database and log growth and reduced server performance can occur over time. For servers at V7.1.7 or later, the Blueprint configuration script enables online database table and index reorganization for most tables by setting the **ALLOWREORGTABLE** and **ALLOWREORGINDEX** server options to YES. Table reorganization is disabled for some larger tables by specifying the **DISABLEREORGTABLE** server option. For the tables in the following list, you can run offline reorganization by using the [Procedure](#):

- BF_AGGREGATED_BITFILES
- BF_BITFILE_EXTENTS
- ARCHIVE_OBJECTS
- BACKUP_OBJECTS

Restriction: Run offline reorganization for the BF_BITFILE_EXTENTS table only if your system includes one or more primary storage pools that were converted to directory-container storage pools.

To run offline reorganization, you must have a file system with enough temporary space to hold an entire table during reorganization. Space within the file systems that are used for database backups can be freed for this purpose.

Because the IBM Spectrum Protect server database grows over time, there might be insufficient space in the database backup file systems to use as free space for the reorganization process. To release space in database backup file systems, you can remove old backup versions.

Complete the following steps to prepare temporary space in the database file systems, and then run offline reorganization.

Procedure

1. Remove the oldest database backups.

For example, to remove the two oldest database backups, issue the following command:

```
delete volhistory type=dbb todate=today-4
```

2. Back up the current version of the database with the **BACKUP DB** command:

```
backup db devc=DBBACK_FILEDEV type=full numstreams=3
```

3. Locate the database backup file system with the most free space to use for the reorganization.
4. Complete the procedure for offline table reorganization. During this step, you might be prompted to back up the database but it is unnecessary for you to do so. Follow the instructions in [technote 1683633](#).

Chapter 7. Next steps

After you complete the setup and configuration for your IBM Spectrum Protect implementation, you can monitor your system and plan for maintenance.

Monitor your system with the IBM Spectrum Protect Operations Center

For more information about the Operations Center, see the following topics.

Getting started with the Operations Center

[Installing and upgrading the Operations Center](#)

Monitoring with the Operations Center

[Monitoring storage solutions](#)

Access the administrative command-line client

The administrative command-line client is installed when you set up your system to run the IBM Spectrum Protect Blueprint configuration script. You can use the administrative client to issue server commands.

For more information about using the **DSMADMC** command to start and stop the administrative client, see [Issuing commands from the administrative client](#).

Review documentation

For documentation in IBM Knowledge Center, see the following links.

Tip: If available, you can display different versions of the same topic in IBM Knowledge Center by using the versions menu at the top of the page.

IBM Spectrum Protect server and client software

- [V8.1.16 documentation](#)

IBM FlashSystem 5000 disk storage systems

[IBM FlashSystem 5000 welcome page](#)

IBM Elastic Storage System

- [IBM Elastic Storage System](#)
- [IBM Spectrum Scale](#)

Optional: Set up data replication by using replication storage rules and subrules

Two or optionally three IBM Spectrum Protect servers that are configured by using the blueprint configuration script can be updated to run replication storage rules. You must specify the parameter **ACTIVE=Yes** to enable the processing of the replication storage rule at the scheduled time.

Before you begin

1. If you are not familiar with the concepts of data replication, review the following information:

Data replication

You can use replication storage rules to create additional copies of data on another server. To learn the basic concepts of data replication, see [Replicating client data to multiple servers](#) in IBM Documentation.

2. Consider whether replication will run in one direction from a source replication server to target replication servers, or if each server will replicate to the other server (acting as both a source and a target replication server). The Blueprint configuration script creates an inactive administrative

schedule for replication with **ACTive=NO** parameter value on all servers. Before this schedule can be activated, you need to complete the replication setup steps below.

3. To optimize data replication operations, ensure that the source replication server and target replication servers have the same hardware configuration, for example:
 - Allocate the same amount of storage capacity on both servers for the database, logs, and storage pools.
 - Use the same type of disks for the database and active log. For example, use solid-state disks for both the database and active log on both servers.
 - Ensure that both servers have the same number of processor cores and a similar amount of read-only memory (RAM). If both servers are used for client backup operations, allocate the same number of processor cores to both servers. However, if the target server is used only for replication, but not for client backup operations, you can allocate half as many processor cores (but no fewer than six) to the target server.

About this task

You can set up data replication by using the **Add Server Pair** wizard in the Operations Center or by following the [Procedure](#).

Procedure

The following manual example assumes that two servers, TAPSRV01 and TAPSRV02, were configured by using the blueprint specifications. The placeholders noted for passwords must match the value that was provided for the server password during the initial configuration. This procedure sets up the data replication so that client nodes' data is backed up to TAPSRV01 and this data is replicated to TAPSRV02.

These steps configure a single storage pool that is used for holding both backup data and replicated data. You can also configure separate storage pools for backup data and replicated data.

1. Set up server-to-server communication.

On TAPSRV01, issue the following command:

```
define server tapsrv02 serverpassword=<secretpassword> hla=tapsrv02.yourdomain.com  
lla=1500
```

On TAPSRV02, issue the following command:

```
define server tapsrv01 serverpassword=<secretpassword> hla=tapsrv01.yourdomain.com  
lla=1500
```

2. Test the communication path.

On TAPSRV01, issue the following command:

```
ping server tapsrv02
```

On TAPSRV02, issue the following command:

```
ping server tapsrv01
```

If the test is successful, you see results similar to the following example:

```
ANR1706I Ping for server 'TAPSRV02' was able to establish a connection.
```

3. Export policy definitions from TAPSRV01 to TAPSRV02. Issue the following command on TAPSRV01:

```
export policy * toserver=tapsrv02
```

4. Define TAPSRV02 as the replication target of TAPSRV01. Issue the following command on TAPSRV01:

```
set replserver tapsrv02
```

5. Enable replication for certain nodes or all nodes. To enable replication for all nodes, issue the following command on TAPSRV01:

```
update node * replstate=enabled
```

6. Define a storage rule to replicate data to the target replication server, TAPSRV02. To define the replication storage rule, REPLRULE1, issue the following command on TAPSRV01:

```
define stgrule replrule1 tapsrv02 actiontype=replicate
```

7. Define an exception to the storage rule, REPLRULE1 to prevent replication of NODE1 by defining a replication subrule. To define the replication subrule, REPLSUBRULE1, issue the following command on TAPSRV01:

```
define subrule replrule1 replsubrule1 node1 actiontype=noreplicating
```

Note: You can replicate data from a source replication server to multiple target replication servers. You must define multiple replication storage rules to configure different target replication servers. Follow the instruction in [step 6](#) to define a replication storage rule for the respective target replication server. If required, follow the instruction in [step 7](#) to define subrules to add exceptions for the respective replication storage rules.

8. On each source replication server, activate the administrative schedule that the Blueprint configuration script created to run replication every day. Issue the following command:

```
update schedule REPLICATE type=admin active=yes
```

Restriction: Ensure that you complete this step only on source replication servers. However, if you are replicating nodes in both directions, and each server is a source and a target replication server, activate the schedule on both servers.

What to do next

To recover data after a disaster, follow the instructions in [Repairing and recovering data in directory-container storage pools](#).

Appendix A. Performance results

You can compare IBM system performance results against your IBM Spectrum Protect storage configuration as a reference for expected performance.

Observed results are based on measurements that were taken in a test lab environment. Test systems were configured according to the Blueprints in this document. Backup-archive clients communicated across a 10 Gb Ethernet connection to the IBM Spectrum Protect server, and deduplicated data was stored in directory-container storage pools. Because many variables can influence throughput in a system configuration, do not expect to see exact matches with the results. Storage pool compression was included in the test configuration on which these performance results are based.

The following typical factors can cause variations in actual performance:

- Average object size of your workload
- Number of client sessions that are used in your environment
- Amount of duplicate data

This information is provided to serve only as a reference.

Extra Small system performance measurements

Data was recorded for a extra small system in the IBM test lab environment.

Metric	Limit or range	Notes
Maximum supported client sessions	75	
Daily amount of new data (before data deduplication)	Up to 1 TB per day ¹	The daily amount of data is how much new data is backed up each day.
Backup ingestion rate	Server-side inline data deduplication	0.5 TB per hour
	Client-side data deduplication	0.75 TB per hour
¹ The daily amount of new data is a range. For more information, see Chapter 2, “Implementation requirements,” on page 3.		

Metric	Range	Notes
Total managed data (size before data deduplication)	10 TB - 40 TB	Total managed data is the volume of data that the server manages, including all versions.

Table 23. Data restore processes

Metric	Number of restore processes	Limit
Throughput of restore processes	1	320 GB per hour
	2	634.5 GB per hour
	4	1116.1 GB per hour
	8	1277.4 GB per hour

Small system performance measurements

Data was recorded for a small system in the IBM test lab environment. These results are taken from the previous blueprint using the FlashSystem 5010.

Table 24. Data intake processes

Metric	Limit or range	Notes
Maximum supported client sessions	250	
Daily amount of new data (before data deduplication)	Up to 10 TB per day ¹	The daily amount of data is how much new data is backed up each day.
Backup ingestion rate	Server-side inline data deduplication	2.9 TB per hour
	Client-side data deduplication	4 TB per hour
¹ The daily amount of new data is a range. For more information, see Chapter 2, “Implementation requirements,” on page 3.		

Table 25. Protected data

Metric	Range	Notes
Total managed data (size before data deduplication)	60 TB - 240 TB	Total managed data is the volume of data that the server manages, including all versions.

Metric	Number of restore processes	Limit
Throughput of restore processes	1	481 GB per hour
	2	716.2 GB per hour
	4	1180.7 GB per hour
	6	1578.1 GB per hour
	8	2250.2 GB per hour

Medium system performance measurements

Data was recorded for a medium system in the IBM test lab environment. The measurements were taken on the previous FlashSystem 5030 configuration and will be updated with results from the FlashSystem 5035 in a future update.

Metric	Limit or range	Notes [®]
Maximum supported client sessions	500	
Daily amount of new data (before data deduplication)	10 - 30 TB per day ¹	The daily amount of data is how much new data is backed up each day.
Backup ingestion rate	Server-side inline data deduplication	5.5 TB per hour
	Client-side data deduplication	8 TB per hour

¹ The daily amount of new data is a range. For more information, see [Chapter 2, "Implementation requirements,"](#) on page 3.

Metric	Range	Notes
Total managed data (size before data deduplication)	360 TB - 1440 TB	Total managed data is the volume of data that the server manages, including all versions.

Metric	Number of restore processes	Limit
Throughput of restore processes	1	456.3 GB per hour
	2	982.54 GB per hour
	4	1367.08 GB per hour
	6	2167.88 GB per hour
	8	2744.43 GB per hour
	10	3335.02 GB per hour
	12	3677.57 GB per hour

Large system performance measurements

Data was recorded for a large system in the IBM test lab environment.

Metric	Limit or range	Notes
Maximum supported client sessions	1000	
Daily amount of new data (before data deduplication)	30 - 100 TB per day ¹	The daily amount of data is how much new data is backed up each day.
Backup ingestion rate	Server-side inline data deduplication	18.1 TB per hour
	Client-side data deduplication	18.4 TB per hour

¹ The daily amount of new data is a range. For more information, see [Chapter 2, "Implementation requirements,"](#) on page 3.

Metric	Range	Notes
Total managed data (size before data deduplication)	1000 TB - 4000 TB	Total managed data is the volume of data that the server manages, including all versions.

Table 32. Data movement

Metric	Number of restore processes	Limit
Throughput of restore processes	1	813.4 GB per hour
	2	1321.0 GB per hour
	4	2546.9 GB per hour
	6	4227.5 GB per hour
	8	5525.7 GB per hour
	10	7816.6 GB per hour
	20	9725.8 GB per hour
	40	10839.6 GB per hour

Workload simulation tool results

Sample data from the workload simulation tool is provided for blueprint test lab systems. Both a storage pool workload and a database workload were tested on each system.

Test results on an IBM Elastic Storage System model GL6 are not currently available.

Extra Small system - storage pool workload

The storage pool workload test included two file systems. The following command was issued:

```
perl tsmdiskperf.pl workload=stgpool fslist=/tsminst1/TSMfile00,/tsminst1/TSMfile01 size=xsmall
```

These results were included in the output:

```
: Average R Throughput (KB/sec):      290214.58
: Average W Throughput (KB/sec):      267109.23
: Avg Combined Throughput (MB/sec):    544.26
: Max Combined Throughput (MB/sec):    716.62
:
: Average IOPS:                        2177.12
: Peak IOPS:                           2866.48 at 08/04/2020 08:59:22
:
: Total elapsed time (seconds):        66
```

Extra Small system - database workload

The database workload test included one file system. The following command was issued:

```
perl tsmdiskperf.pl workload=db fslist=/tsminst1/TSMdbspace00 size=xsmall
```

These results were included in the output:

```
: Average R Throughput (KB/sec):      737.70
: Average W Throughput (KB/sec):      869.19
: Avg Combined Throughput (MB/sec):    1.57
: Max Combined Throughput (MB/sec):    4.72
:
: Average IOPS:                        174.37
: Peak IOPS:                           451.86 at 08/04/202 09:34:17
```

```
:  
: Total elapsed time (seconds):          47
```

Small system - storage pool workload

The storage pool workload test included eight file systems. The following command was issued:

```
perl tsmdiskperf.pl workload=stgpool fslist=/tsminst1/TSMfile00,  
/tsminst1/TSMfile01,/tsminst1/TSMfile02,/tsminst1/TSMfile03,  
/tsminst1/TSMfile04,/tsminst1/TSMfile05,  
/tsminst1/TSMfile06,/tsminst1/TSMfile07
```

These results were included in the output:

```
: Average R Throughput (KB/sec):          195446.67  
: Average W Throughput (KB/sec):          196250.62  
: Avg Combined Throughput (MB/sec):       382.52  
: Max Combined Throughput (MB/sec):       585.72  
:  
: Average IOPS:                           3060.89  
: Peak IOPS:                               4686.57 at 04/16/2019 02:45:40  
:  
: Total elapsed time (seconds):           399
```

Small system - database workload

The database workload test included four file systems. The following command was issued:

```
perl tsmdiskperf.pl workload=db fslist=/tsminst1/TSMdbspace00,  
/tsminst1/TSMdbspace01,/tsminst1/TSMdbspace02,/tsminst1/TSMdbspace03
```

These results were included in the output:

```
: Average R Throughput (KB/sec):          22811.60  
: Average W Throughput (KB/sec):          22732.55  
: Avg Combined Throughput (MB/sec):       44.48  
: Max Combined Throughput (MB/sec):       71.73  
:  
: Average IOPS:                           5693.30  
: Peak IOPS:                               9182.15 at 04/16/2019 12:40:33  
:  
: Total elapsed time (seconds):           699
```

Medium system - storage pool workload

The storage pool workload test included 10 file systems. The following command was issued:

```
perl tsmdiskperf.pl workload=stgpool  
fslist=/tsminst1/TSMfile00,/tsminst1/TSMfile01,  
/tsminst1/TSMfile02,/tsminst1/TSMfile03,  
/tsminst1/TSMfile04,/tsminst1/TSMfile05,  
/tsminst1/TSMfile06,/tsminst1/TSMfile07,  
/tsminst1/TSMfile08,/tsminst1/TSMfile09
```

These results were included in the output:

```
: Average R Throughput (KB/sec):          995464.88  
: Average W Throughput (KB/sec):          996739.71  
: Avg Combined Throughput (MB/sec):       1945.51  
: Max Combined Throughput (MB/sec):       2778.50  
:  
: Average IOPS:                           10915.62  
: Peak IOPS:                               15774.13 at 06/05/2018 10:42:56  
:  
: Total elapsed time (seconds):           196
```

Medium system - database workload

The database workload test included eight file systems. The following command was issued:

```
perl tsmdiskperf.pl workload=db
fslist=/tsminst1/TSMdbspace00,/tsminst1/TSMdbspace01,/tsminst1/TSMdbspace02,
/tsminst1/TSMdbspace03,/tsminst1/TSMdbspace04,/tsminst1/TSMdbspace05,
/tsminst1/TSMdbspace06,/tsminst1/TSMdbspace07
```

These results were included in the output:

```
: Average R Throughput (KB/sec):      49052.01
: Average W Throughput (KB/sec):      48982.18
: Avg Combined Throughput (MB/sec):    95.74
: Max Combined Throughput (MB/sec):    155.94
:
: Average IOPS:                        12238.62
: Peak IOPS:                           19956.72 at 06/05/2018 03:52:08
:
: Total elapsed time (seconds):        651
```

Large IBM FlashSystem system - storage pool workload

The storage pool workload test included 32 file systems. The following command was issued:

```
perl tsmdiskperf.pl workload=stgpool fslist=/tsminst1/TSMfile00,
/tsminst1/TSMfile01,/tsminst1/TSMfile02,/tsminst1/TSMfile03,
/tsminst1/TSMfile04,/tsminst1/TSMfile05,/tsminst1/TSMfile06,
/tsminst1/TSMfile07,/tsminst1/TSMfile08,/tsminst1/TSMfile09,
/tsminst1/TSMfile10,/tsminst1/TSMfile11,/tsminst1/TSMfile12,
/tsminst1/TSMfile13,/tsminst1/TSMfile14,/tsminst1/TSMfile15,
/tsminst1/TSMfile16,/tsminst1/TSMfile17,/tsminst1/TSMfile18,
/tsminst1/TSMfile19,/tsminst1/TSMfile20,/tsminst1/TSMfile21,
/tsminst1/TSMfile22,/tsminst1/TSMfile23,/tsminst1/TSMfile24,
/tsminst1/TSMfile25,/tsminst1/TSMfile26,/tsminst1/TSMfile27,
/tsminst1/TSMfile28,/tsminst1/TSMfile29,/tsminst1/TSMfile30,
/tsminst1/TSMfile31
```

These results were included in the output:

```
: Average R Throughput (KB/sec):      2135345.73
: Average W Throughput (KB/sec):      2132259.14
: Avg Combined Throughput (MB/sec):    4167.58
: Max Combined Throughput (MB/sec):    5468.96
:
: Average IOPS:                        16673.23
: Peak IOPS:                           21875.84 at 09/07/2022 03:21:10
:
: Total elapsed time (seconds):        147
```

Large IBM FlashSystem system - database workload

The database workload test included 12 file systems. The following command was issued:

```
perl /code/tsmconfig/tsmdiskperf.pl workload=db
fslist=/tsminst1/TSMdbspace00,/tsminst1/TSMdbspace01,
/tsminst1/TSMdbspace02,/tsminst1/TSMdbspace03,
/tsminst1/TSMdbspace04,/tsminst1/TSMdbspace05,
/tsminst1/TSMdbspace06,/tsminst1/TSMdbspace07,
/tsminst1/TSMdbspace08,/tsminst1/TSMdbspace09,
/tsminst1/TSMdbspace10,/tsminst1/TSMdbspace11
```

These results were included in the output:

```
: Average R Throughput (KB/sec):      279037.50
: Average W Throughput (KB/sec):      291834.51
: Avg Combined Throughput (MB/sec):    557.49
: Max Combined Throughput (MB/sec):    844.46
:
: Average IOPS:                        69834.90
```

```
: Peak IOPS: 108091.16 at 09/07/2022 02:56:29
:
: Total elapsed time (seconds): 116
```

Large IBM Elastic Storage System system - storage pool workload

The storage pool workload test included execution from a single server using 64 directories within one file system.

```
perl tsmdiskperf.pl workload=stgpool fslist=/gpfs/tsmp06_tsmppool/tsmdiskperf/data/spg1-INST01/
test01,/gpfs/tsmp06_tsmppool/tsmdiskperf/data/spg1-INST01/test02,
.../gpfs/tsmp06_tsmppool/tsmdiskperf/data/spg1-INST01/test64
```

These results were included in the output:

```
: Average R Throughput (KB/sec): 7507649.89
: Average W Throughput (KB/sec): 7491784.55
: Avg Combined Throughput (MB/sec): 14647.89
: Max Combined Throughput (MB/sec): 17489.01
:
: Average IOPS: 57340.97
: Peak IOPS: 68414.25 at 1626266586
:
: Total elapsed time (seconds): 80
```

When running the 64 directory storage pool workload from four servers simultaneously, the total average combined throughput from all servers exceeded 35,000 MB/sec.

Large IBM Elastic Storage System - database workload

The database workload test included execution from a single server using 12 directories within one file system from the ESS 3000 flash storage.

```
perl tsmdiskperf.pl workload=stgpool fslist=/gpfs/tsmp06_tsmdb/tsmdiskperf/data/spg1-INST01/
test01,/gpfs/tsmp06_tsmdb/tsmdiskperf/data/spg1-INST01/test02,
.../gpfs/tsmp06_tsmdb/tsmdiskperf/data/spg1-INST01/test12
```

These results were included in the output:

```
: Average R Throughput (KB/sec): 181022.19
: Average W Throughput (KB/sec): 180649.21
: Avg Combined Throughput (MB/sec): 353.19
: Max Combined Throughput (MB/sec): 640.81
:
: Average IOPS: 44221.35
: Peak IOPS: 80289.58 at 1626267223
:
: Total elapsed time (seconds): 186
```

When running the 12 directory database workload from four servers simultaneously, the total combined average IOPS from all servers exceeded 130,000 IOPS.

Appendix B. Configuring the disk system by using commands

You can use the IBM FlashSystem command line to configure storage arrays and volumes on the disk system. Example procedures are provided for the 5015 (small), 5035 (medium), and 5200 (large) systems.

Refer to [Chapter 3, “Storage configuration blueprints,”](#) on page 15 for layout specifications.

Small system

1. Connect to and log in to the disk system by issuing the **ssh** command. For example:

```
ssh superuser@your5010hostname
```

2. List drive IDs for each type of disk so that you can create the managed disk (MDisk) arrays in Step “4” on page 65. Issue the **lsdrive** command. The output can vary, based on slot placement for the different disks. The output is similar to the following example:

id	status	use	tech_type	capacity	...	enclosure_id	slot_id	...
0	online	candidate	tier0_flash	1.45TB		1	3	
1	online	candidate	tier0_flash	1.45TB		1	4	
2	online	candidate	tier0_flash	1.45TB		1	1	
3	online	candidate	tier0_flash	1.45TB		1	2	
4	online	candidate	tier0_flash	1.45TB		1	5	
5	online	candidate	tier_nearline	3.6TB		2	6	
6	online	candidate	tier_nearline	3.6TB		2	1	
7	online	candidate	tier_nearline	3.6TB		2	7	
8	online	candidate	tier_nearline	3.6TB		2	10	
9	online	candidate	tier_nearline	3.6TB		2	5	
10	online	candidate	tier_nearline	3.6TB		2	4	
11	online	candidate	tier_nearline	3.6TB		2	2	
12	online	candidate	tier_nearline	3.6TB		2	9	
13	online	candidate	tier_nearline	3.6TB		2	11	
14	online	candidate	tier_nearline	3.6TB		2	3	
15	online	candidate	tier_nearline	3.6TB		2	12	
16	online	candidate	tier_nearline	3.6TB		2	8	
17	online	candidate	tier_nearline	3.6TB		3	6	
18	online	candidate	tier_nearline	3.6TB		3	12	
19	online	candidate	tier_nearline	3.6TB		3	9	
20	online	candidate	tier_nearline	3.6TB		3	4	
21	online	candidate	tier_nearline	3.6TB		3	11	
22	online	candidate	tier_nearline	3.6TB		3	5	
23	online	candidate	tier_nearline	3.6TB		3	2	
24	online	candidate	tier_nearline	3.6TB		3	10	
25	online	candidate	tier_nearline	3.6TB		3	8	
26	online	candidate	tier_nearline	3.6TB		3	1	
27	online	candidate	tier_nearline	3.6TB		3	7	
28	online	candidate	tier_nearline	3.6TB		3	3	

3. Create the MDisk groups for the IBM Spectrum Protect database and storage pool. Issue the **mkmdiskgroup** command for each pool, specifying 256 for the extent size:

```
mkmdiskgrp -name db_grp0 -ext 256  
mkmdiskgrp -name stgpool_grp0 -ext 256
```

4. Create MDisk arrays by using **mkdistributedarray** commands. Specify the commands to add the MDisk arrays to the data pools that you created in the previous step. For example:

```
mkdistributedarray -name db_array0 -level raid5 -driveclass 2 -drivecount 4 -stripewidth 3  
-rebuildareas 1 -strip 256 db_grp0  
mkdistributedarray -name stgpool_array0 -level raid6 -driveclass 1 -drivecount 24  
-stripewidth 12 -rebuildareas 1 -strip 256 stgpool_grp0
```

5. Create the storage volumes for the system. Issue the **mkvdisk** command for each volume, specifying the volume sizes in MB. For example:

```
mkvdisk -mdiskgrp db_grp0 -size 343296 -name db_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 343296 -name db_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 343296 -name db_02 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 343296 -name db_03 -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp db_grp0 -size 148736 -name alog -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 1244928 -name archlog -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 3303398 -unit mb -name backup_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 3303398 -unit mb -name backup_01 -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 15859710 -unit mb -name filepool_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 15859710 -unit mb -name filepool_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 15859710 -unit mb -name filepool_02 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 15859710 -unit mb -name filepool_03 -iogrp 0 -nofmtdisk
```

6. Create a logical host object by using the **mkhost** command. Specify the Fibre Channel WWPNs from your operating system and specify the name of your host. To obtain the WWPNs from your system, follow the instructions in “Step 1: Set up and configure hardware” on page 23.

For example, to create a host that is named *hostone* with a list that contains FC WWPNs *1000090FA3D8F12* and *1000090FA49009E*, issue the following command:

```
mkhost -name hostone -fcwwpn 1000090FA3D8F12:1000090FA49009E -iogrp 0
-type=generic -force
```

7. Map the volumes that you created in Step “5” on page 66 to the new host. Issue the **mkvdiskhostmap** command for each volume. For example, issue the following commands where *hostname* is the name of your host:

```
mkvdiskhostmap -host hostname -scsi 0 db_00
mkvdiskhostmap -host hostname -scsi 1 db_01
mkvdiskhostmap -host hostname -scsi 2 db_02
mkvdiskhostmap -host hostname -scsi 3 db_03

mkvdiskhostmap -host hostname -scsi 4 alog

mkvdiskhostmap -host hostname -scsi 5 archlog

mkvdiskhostmap -host hostname -scsi 6 backup_0
mkvdiskhostmap -host hostname -scsi 7 backup_1

mkvdiskhostmap -host hostname -scsi 8 filepool_00
mkvdiskhostmap -host hostname -scsi 9 filepool_01
mkvdiskhostmap -host hostname -scsi 10 filepool_02
mkvdiskhostmap -host hostname -scsi 11 filepool_03
```

Medium system

1. Connect to and log in to the disk system by issuing the **ssh** command. For example:

```
ssh superuser@your5010hostname
```

2. Increase the memory that is available for the RAID5s to 125 MB by issuing the **chiogrp** command:

```
chiogrp -feature raid -size 125 io_grp0
```

3. List drive IDs for each type of disk so that you can create the MDisk arrays in Step “5” on page 67. Issue the **lsdrive** command. The output can vary, based on slot placement for the different disks. The output is similar to the following example:

```
IBM_Storwize:tapv5kg:superuser>lsdrive
id status use tech_type capacity enclosure_id slot_id drive_class_id
0 online member tier_nearline 5.5TB 1 26 0
1 online member tier_nearline 5.5TB 1 44 0
2 online member tier_nearline 5.5TB 1 1 0
3 online member tier_nearline 5.5TB 1 34 0
```

4	online	member	tier_nearline	5.5TB	1	20	0
5	online	member	tier_nearline	5.5TB	1	25	0
< ... >							
91	online	member	tier_nearline	5.5TB	1	2	0
92	online	member	tier1_flash	1.7TB	2	4	1
93	online	member	tier1_flash	1.7TB	2	1	1
94	online	member	tier1_flash	1.7TB	2	3	1
95	online	member	tier1_flash	1.7TB	2	6	1
96	online	member	tier1_flash	1.7TB	2	5	1
97	online	member	tier1_flash	1.7TB	2	2	1

4. Create the MDisk groups for the IBM Spectrum Protect database and storage pool. Issue the **mkmdiskgroup** command for each pool, specifying 1024 for the extent size:

```
mkmdiskgrp -name db_grp0 -ext 1024
mkmdiskgrp -name stgpool_grp0 -ext 1024
```

5. Create MDisk arrays by using **mkdistributedarray** commands. Specify the commands to add the MDisk arrays to the data pools that you created in the previous step.

For example:

```
mkdistributedarray -name db_array0 -level raid6 -driveclass 1 -drivecount 6 -stripewidth 5
-rebuildareas 1 -strip 256 db_grp0
mkdistributedarray -name stgpool_array0 -level raid6 -driveclass 0 -drivecount 46
-stripewidth 12 -rebuildareas 2 -strip 256 stgpool_grp0
mkdistributedarray -name stgpool_array1 -level raid6 -driveclass 0 -drivecount 46
-stripewidth 12 -rebuildareas 2 -strip 256 stgpool_grp0
```

6. Create the storage volumes for the system. Issue the **mkvdisk** command for each volume, specifying the volume sizes in MB. For example:

```
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_02 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_03 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_04 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_05 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_06 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 656999 -name db_07 -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp db_grp0 -size 150528 -name alog -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 2097152 -name archlog -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 15728640 -name backup_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 15728640 -name backup_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 15728640 -name backup_02 -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_02 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_03 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_04 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_05 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_06 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_07 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_08 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_09 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_10 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 30648320 -unit mb -name filepool_11 -iogrp 0 -nofmtdisk
```

7. Create a logical host object by using the **mkhost** command. Specify the Fibre Channel WWPNs from your operating system and specify the name of your host. To obtain the WWPNs from your system, follow the instructions in [“Step 1: Set up and configure hardware”](#) on page 23.

For example, to create a host that is named *hostone* with a list that contains FC WWPNs *10000090FA3D8F12* and *10000090FA49009E*, issue the following command:

```
mkhost -name hostone -fcwwpn 10000090FA3D8F12:10000090FA49009E -iogrp 0
-type=generic -force
```

- Map the volumes that you created in Step “6” on page 67 to the new host. Issue the **mkvdiskhostmap** command for each volume. For example, issue the following commands where *hostname* is the name of your host:

```

mkvdiskhostmap -host hostname -scsi 0 db_00
mkvdiskhostmap -host hostname -scsi 1 db_01
mkvdiskhostmap -host hostname -scsi 2 db_02
mkvdiskhostmap -host hostname -scsi 3 db_03
mkvdiskhostmap -host hostname -scsi 4 db_04
mkvdiskhostmap -host hostname -scsi 5 db_05
mkvdiskhostmap -host hostname -scsi 6 db_06
mkvdiskhostmap -host hostname -scsi 7 db_07

mkvdiskhostmap -host hostname -scsi 8 alog

mkvdiskhostmap -host hostname -scsi 9 archlog

mkvdiskhostmap -host hostname -scsi 10 backup_00
mkvdiskhostmap -host hostname -scsi 11 backup_01
mkvdiskhostmap -host hostname -scsi 12 backup_02

mkvdiskhostmap -host hostname -scsi 13 filepool_00
mkvdiskhostmap -host hostname -scsi 14 filepool_01
mkvdiskhostmap -host hostname -scsi 15 filepool_02
mkvdiskhostmap -host hostname -scsi 16 filepool_03
mkvdiskhostmap -host hostname -scsi 17 filepool_04
mkvdiskhostmap -host hostname -scsi 18 filepool_05
mkvdiskhostmap -host hostname -scsi 19 filepool_06
mkvdiskhostmap -host hostname -scsi 20 filepool_07
mkvdiskhostmap -host hostname -scsi 21 filepool_08
mkvdiskhostmap -host hostname -scsi 22 filepool_09
mkvdiskhostmap -host hostname -scsi 23 filepool_10
mkvdiskhostmap -host hostname -scsi 24 filepool_11

```

Large system

- Connect to and log in to the disk system by issuing the **ssh** command. For example:

```
ssh superuser@your5200hostname
```

- Increase the memory that is available for the RAIDs to 125 MB by issuing the **chiogrp** command:

```
chiogrp -feature raid -size 125 io_grp0
```

- List drive IDs for each type of disk so that you can create the MDisk arrays in Step “5” on page 68. Issue the **lsdrive** command. The output can vary, based on slot placement for the different disks. The output is similar to what is returned for small and medium systems.
- Create the MDisk groups for the IBM Spectrum Protect database and storage pool. Issue the **mkmdiskgrp** command for each pool, specifying 1024 for the extent size:

```
mkmdiskgrp -name db_grp0 -ext 1024
mkmdiskgrp -name stgpool_grp0 -ext 1024
```

- Create arrays by using the **mkdistributedarray** command. Specify the commands to add the MDisk arrays to the data pools that you created in the previous step.

For example:

```

mkdistributedarray -name db_array0 -level raid6 -driveclass 0 -drivecount 9 -stripewidth 8
-rebuildareas 1 -strip 256 db_grp0
mkdistributedarray -name stgpool_array0 -level raid6 -driveclass 1 -drivecount 46
-stripewidth 12 -rebuildareas 2 -strip 256 stgpool_grp0
mkdistributedarray -name stgpool_array1 -level raid6 -driveclass 1 -drivecount 46
-stripewidth 12 -rebuildareas 2 -strip 256 stgpool_grp0
mkdistributedarray -name stgpool_array2 -level raid6 -driveclass 1 -drivecount 46
-stripewidth 12 -rebuildareas 2 -strip 256 stgpool_grp0
mkdistributedarray -name stgpool_array3 -level raid6 -driveclass 1 -drivecount 46
-stripewidth 12 -rebuildareas 2 -strip 256 stgpool_grp0

```

- Create the storage volumes for the system. Issue the **mkvdisk** command for each volume, specifying the volume sizes in MB.

For example:

```
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_02 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_03 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_04 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_05 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_06 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_07 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_08 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_09 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_10 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp db_grp0 -size 858000 -unit mb -name db_11 -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp db_grp0 -size 563200 -unit mb -name alog -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 4200000 -unit mb -name archlog -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 18874368 -unit mb -name backup_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 18874368 -unit mb -name backup_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 18874368 -unit mb -name backup_02 -iogrp 0 -nofmtdisk

mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_00 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_01 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_02 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_03 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_04 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_05 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_06 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_07 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_08 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_09 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_10 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_11 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_12 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_13 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_14 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_15 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_16 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_17 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_18 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_19 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_20 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_21 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_22 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_23 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_24 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_25 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_26 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_27 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_28 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_29 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_30 -iogrp 0 -nofmtdisk
mkvdisk -mdiskgrp stgpool_grp0 -size 32856064 -unit mb -name filepool_31 -iogrp 0 -nofmtdisk
```

7. Create a logical host object by using the **mkhost** command. Specify the Fibre Channel WWPNs from your operating system and specify the name of your host. For instructions about obtaining the WWPNs from your system, see “Step 1: Set up and configure hardware” on page 23.

For example, to create a host that is named *hostone* with a list that contains FC WWPNs *10000090FA3D8F12* and *10000090FA49009E*, issue the following command:

```
mkhost -name hostone -fcwwpn
10000090FA3D8F12:10000090FA3D8F13:10000090FA49009E:10000090FA49009F -iogrp 0
-type=generic -force
```

8. Map the volumes that you created in Step “6” on page 68 to the new host. Issue the **mkvdiskhostmap** command for each volume. For example, issue the following commands where *hostname* is the name of your host:

```
mkvdiskhostmap -host hostname -scsi 0 db_00
mkvdiskhostmap -host hostname -scsi 1 db_01
mkvdiskhostmap -host hostname -scsi 2 db_02
mkvdiskhostmap -host hostname -scsi 3 db_03
mkvdiskhostmap -host hostname -scsi 4 db_04
mkvdiskhostmap -host hostname -scsi 5 db_05
```

```
mkvdiskhostmap -host hostname -scsi 6 db_06
mkvdiskhostmap -host hostname -scsi 7 db_07
mkvdiskhostmap -host hostname -scsi 8 db_08
mkvdiskhostmap -host hostname -scsi 9 db_09
mkvdiskhostmap -host hostname -scsi 10 db_10
mkvdiskhostmap -host hostname -scsi 11 db_11

mkvdiskhostmap -host hostname -scsi 12 alog

mkvdiskhostmap -host hostname -scsi 13 archlog

mkvdiskhostmap -host hostname -scsi 14 backup_00
mkvdiskhostmap -host hostname -scsi 15 backup_01
mkvdiskhostmap -host hostname -scsi 16 backup_02

mkvdiskhostmap -host hostname -scsi 17 filepool_00
mkvdiskhostmap -host hostname -scsi 18 filepool_01
mkvdiskhostmap -host hostname -scsi 19 filepool_02
mkvdiskhostmap -host hostname -scsi 20 filepool_03
mkvdiskhostmap -host hostname -scsi 21 filepool_04
mkvdiskhostmap -host hostname -scsi 22 filepool_05
mkvdiskhostmap -host hostname -scsi 23 filepool_06
mkvdiskhostmap -host hostname -scsi 24 filepool_07
mkvdiskhostmap -host hostname -scsi 25 filepool_08
mkvdiskhostmap -host hostname -scsi 26 filepool_09
mkvdiskhostmap -host hostname -scsi 27 filepool_10
mkvdiskhostmap -host hostname -scsi 28 filepool_11
mkvdiskhostmap -host hostname -scsi 29 filepool_12
mkvdiskhostmap -host hostname -scsi 30 filepool_13
mkvdiskhostmap -host hostname -scsi 31 filepool_14
mkvdiskhostmap -host hostname -scsi 32 filepool_15
mkvdiskhostmap -host hostname -scsi 33 filepool_16
mkvdiskhostmap -host hostname -scsi 34 filepool_17
mkvdiskhostmap -host hostname -scsi 35 filepool_18
mkvdiskhostmap -host hostname -scsi 36 filepool_19
mkvdiskhostmap -host hostname -scsi 37 filepool_20
mkvdiskhostmap -host hostname -scsi 38 filepool_21
mkvdiskhostmap -host hostname -scsi 39 filepool_22
mkvdiskhostmap -host hostname -scsi 40 filepool_23
mkvdiskhostmap -host hostname -scsi 41 filepool_24
mkvdiskhostmap -host hostname -scsi 42 filepool_25
mkvdiskhostmap -host hostname -scsi 43 filepool_26
mkvdiskhostmap -host hostname -scsi 44 filepool_27
mkvdiskhostmap -host hostname -scsi 45 filepool_28
mkvdiskhostmap -host hostname -scsi 46 filepool_29
mkvdiskhostmap -host hostname -scsi 47 filepool_30
mkvdiskhostmap -host hostname -scsi 48 filepool_31
```

Appendix C. Using a response file with the Blueprint configuration script

You can run the Blueprint configuration script in non-interactive mode by using a response file to set your configuration choices.

Three response files are provided with the Blueprint configuration script. If you plan to set up a system by using all default values, you can run the configuration script in non-interactive mode by using one of the following response files:

Small system

responsefilesmall.txt

Medium system

responsefilemed.txt

Large system

- Storwize systems: responsefilelarge.txt
- IBM Elastic Storage System systems: responsefile_ess.txt

The files are pre-filled with default configuration values for the small, medium, and large systems and do not require updates.

If you want to customize your responses for a system, use the following table with your “[Planning worksheets](#)” on [page 10](#) to update one of the default response files. The values that are used in the response file correspond to values that you recorded in the *Your value* column of the worksheet.

Response file value	Corresponding value from the planning worksheet
serverscale	Not recorded in the planning worksheet. Enter a value of <i>S</i> for a small system, <i>M</i> for a medium system, or <i>L</i> for a large system.
db2user	Db2 instance owner ID
db2userpw	Db2 instance owner password
db2group	Primary group for the Db2 instance owner ID
db2userhomedir	Home directory for the Db2 instance owner ID. By default, this directory is created in the /home file system. For IBM Elastic Storage System configurations, the preferred method is to use a directory in the shared IBM Spectrum Scale file system.
instdirmountpoint	Directory for the server instance
db2dirpaths	Directories for the database
tsmstgpaths	Directories for storage
actlogpath	Directory for the active log
archlogpath	Directory for the archive log
dbbackdirpaths	Directories for database backup
backupstarttime	Schedule start time
tsmsysadminid	IBM Spectrum Protect administrator ID
tsmsysadminidpw	IBM Spectrum Protect administrator ID password

Response file value	Corresponding value from the planning worksheet
tcpport	TCP/IP port address for communications with the IBM Spectrum Protect server. Enter the value that you specified for Step “6” on page 27.
servername	Server name
serverpassword	Server password

Appendix D. Using predefined client schedules

The Blueprint configuration script creates several client schedules during server configuration. To use these schedules, you must complete configuration steps on the client system.

Table 33 on page 73 lists the predefined schedules that are created on the server. The schedule names and descriptions are based on the default backup schedule start time of 10 PM. If you changed this start time during server configuration, the predefined client schedules on your system are named according to that start time. Information about updating client schedules to use with the IBM Spectrum Protect server is provided in the sections that follow the table.

For complete information about scheduling client backup operations, see your client documentation.

Client	Schedule name	Schedule description
IBM Spectrum Protect for Databases: Data Protection for Oracle	ORACLE_DAILYFULL_10PM	Oracle Daily FULL backup that starts at 10 PM
IBM Spectrum Protect for Databases: Data Protection for Microsoft SQL Server	SQL_DAILYFULL_10PM	Microsoft SQL Daily FULL backup that starts at 10 PM
IBM Spectrum Protect backup-archive client	FILE_INCRFOREVER_10PM	File incremental-forever backup that starts at 10 PM
IBM Spectrum Protect for Mail: Data Protection for HCL Domino®	DOMINO_DAILYFULL_10PM	Daily FULL backup that starts at 10 PM
IBM Spectrum Protect for Mail: Data Protection for Microsoft Exchange Server	EXCHANGE_DAILYFULL_10PM	FULL backup that starts at 10 PM
IBM Spectrum Protect for Virtual Environments: Data Protection for Microsoft Hyper-V	HYPERV_FULL_10PM	Hyper-V full backup that starts at 10 PM

Data Protection for Oracle

Data Protection for Oracle does not include a sample backup file. You can create a script or .bat command file and update the **OBJECTS** parameter for the predefined schedule by using the **UPDATE SCHEDULE** server command. Specify the full path to the command file on the client system unless you save the command file in the client installation directory. Then, you must provide only the file name.

For example, to update the ORACLE_DAILYFULL_10PM schedule that is in the DATABASE domain, issue the following command. Specify the name of the command file that you want to use in the client installation directory. In this example, the command file is named schedcmdfile.bat.

```
update schedule database oracle_dailyfull_10pm obj=schedcmdfile.bat
```

Data Protection for Microsoft SQL Server

The sample schedule file that is included with Data Protection for Microsoft SQL Server is named sqlfull.cmd. This file can be customized for use with IBM Spectrum Protect server. If you save the file to the client installation directory on the client system, you do not have to update the predefined schedule to include the full path to the file.

Backup-archive client

When you use the predefined schedule for backup-archive clients, the server processes objects as they are defined in the client options file, unless you specify a file to run a command or macro. For information about setting the domain, include, and exclude options for backup operations, see the online product documentation:

- [Client options reference \(V6.4\)](#)
- [Client options reference \(V7.1\)](#)
- [Client options reference \(V8.1\)](#)

Data Protection for HCL Domino

The sample schedule file that is included with Data Protection for HCL Domino is named `domsel.cmd`. This file can be customized for use with IBM Spectrum Protect server. If you save the file to the client installation directory on the client system, you do not have to update the predefined schedule to include the full path to the file.

Data Protection for Microsoft Exchange Server

The sample schedule file that is included with Data Protection for Microsoft Exchange Server is named `excfull.cmd`. This file can be customized for use with IBM Spectrum Protect server. If you save the file to the client installation directory on the client system, you do not have to update the predefined schedule to include the full path to the file.

Data Protection for Microsoft Hyper-V

No sample schedule file is provided with Data Protection for Microsoft Hyper-V. To create a `.cmd` file that can back up multiple virtual machines, complete the following steps:

1. Update the client options file to include the following settings:

```
commethod      tcpip
tcpport        1500
TCPServeraddress <IBM Spectrum Protect server name>
nodename       <node name>
passwordaccess generate
vmbackuptype   hypervfull
```

2. For each virtual machine that you want to back up, create a separate script file. A unique file is needed to ensure that a log is saved for each backup. For example, create a file that is named `hvvm1.cmd`. Include the backup command, the name of the virtual machine, the client options file, and the log file that you want to create on the first line. On the second line, include the word `exit`.

For example:

```
dsmc backup vm "tsmhyp1vm3" -optfile=dsm-hv.opt >> hv_backup_3.log
exit
```

Repeat this step for each virtual machine that you want to back up.

3. Create a backup schedule file, for example, `hv_backup.cmd`.
4. Add an entry to `hv_backup.cmd` for each virtual machine script file that you created. For example:

```
start hvvm1.cmd
choice /T 10 /C X /D X /N > NUL
start hvvm2.cmd
choice /T 10 /C X /D X /N > NUL
start hvvm3.cmd
choice /T 10 /C X /D X /N > NUL
hvvm4.cmd
```

5. Issue the **UPDATE SCHEDULE** server command to update the predefined `HYPERV_FULL_10PM` schedule. Specify the full path for the Hyper-V backup schedule file location in the **OBJECTS** parameter.

IBM Spectrum Protect for Virtual Environments

To create new schedules, use the Data Protection for VMware vCenter plug-in GUI.

Appendix E. Modification of blueprint configurations

If you want to customize the configurations that are detailed in this document, plan carefully.

Consider the following before you deviate from the blueprint specifications:

- If you want to extend the usable storage for your system by adding storage enclosures, you must also add storage for the IBM Spectrum Protect database. Increase the database storage by approximately 1% of the additional total amount of managed data that will be protected (size before data deduplication).
- You can use Linux operating systems other than Red Hat Enterprise Linux, but the following caveats apply:
 - The version and operating system must be supported for use with the IBM Spectrum Protect server.
 - Additional configuration steps or modifications to steps for installation and configuration might be needed.
- If you use other storage systems, performance measurements that are reported for the blueprint configurations are not guaranteed to match your customization.
- In general, no guarantees can be made for a customized environment. Test the environment to ensure that it meets your business requirements.

Appendix F. Troubleshooting

At the time of publication, the following issue was known.

Slow throughput after server installation

In some cases, following a new installation of IBM Spectrum Protect, the server might experience slow throughput. This condition can be caused by a delay in the Db2 runstats operation, which optimizes how queries are performed. An indication of this issue is that the Db2 process db2sysc is using a large amount of CPU processing as compared to the amount of processing that is used by the server.

To resolve this problem, you can start runstats processing manually. Issue the following command from the administrative command-line interface:

```
dsmadm > runstats all
```

Appendix G. 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 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and Web Content Accessibility Guidelines (WCAG) 2.0 (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 Documentation is enabled for 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 \(www.ibm.com/able\)](http://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.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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>.

Index

A

accessibility features [81](#)
ACTIVELOGSIZE server option [44](#)
ALLOWREORGINDEX server option [51](#)
ALLOWREORGTABLE server option [51](#)
ARCHLOGCOMPRESS server option [44](#)

B

BIOS settings [23](#)
blueprint
 customization [77](#)
 large system [20](#), [21](#), [31](#)
 medium system [18](#)
 small system [17](#)
blueprint configuration script
 compression option [41](#)
 configuring with [41](#)
 planning for [10](#)
 response file [71](#)
 testing [48](#)
 troubleshooting [48](#)

C

COMMTIMEOUT server option [44](#)
compression [41](#)
configuration
 clean up script [48](#)
 customizing [77](#)
 prerequisites [5](#)
configuring
 Db2 database [41](#)
 disk systems [23](#)
 file systems [29](#), [30](#)
 hardware [23](#)
 IBM Elastic Storage System [31](#)
 IBM Spectrum Protect server [41](#)
 RAID arrays [15](#), [23](#)
 Storwize V5010 system [65](#)
 Storwize V5030 system [65](#)
 TCP/IP settings [25](#)
 volumes [15](#), [23](#)

D

data replication [53](#)
database
 configuration of [41](#)
 reorganizing tables and indexes [51](#)
Db2 -locklist parameter [43](#)
DEDUPDELETIONTHREADS server option [44](#)
DEDUPREQUIRESBACKUP server option [44](#)
DEFINE ASSOCIATION server command [49](#)
DEVCONFIG server option [44](#)

directories
 creating [29](#), [30](#)
disability [81](#)
DISABLEREORGTABLE server option [51](#)
disk
 configuring with command line [65](#)
documentation [53](#)
dmserv.opt file [44](#)

E

EXPINTERVAL server option [44](#)

F

file systems
 creating [29](#), [30](#)
 planning for [10](#)

H

hardware requirements [5](#), [7–9](#)

I

IBM Elastic Storage System
 configuring [31](#)
 hardware requirements [9](#)
 storage blueprint [21](#)
IBM FlashSystem 5015
 storage configuration for small systems [17](#)
IBM FlashSystem 5035
 storage configuration for medium systems [18](#)
IBM FlashSystem 5200)
 storage configuration for large systems [20](#)
IBM Knowledge Center [53](#)
IBM Spectrum Protect directories
 planning for [10](#)
IBM Spectrum Protect server
 cleaning up after a failed configuration attempt [48](#)
 configuring [41](#)
 schedules [41](#)
IBM Spectrum Scale
 configuring [31](#)
 installing [31](#)
IDLETIMEOUT server option [44](#)
installing
 IBM Spectrum Protect backup-archive client [37](#)
 IBM Spectrum Protect server [38](#), [39](#)
 IBM Spectrum Scale [31](#)
 obtaining IBM Spectrum Protectserver installation packages [38](#)
 operating system [25](#)
 Red Hat Enterprise Linux [25](#)
iostat command [34](#)

K

kernel parameters [41](#)
keyboard [81](#)

L

Ideedee program [34](#)
Lenovo ThinkSystem SR650
 hardware requirements [7–9](#)
Linux commands
 dd command [34](#)
 iostat command [34](#)
 passwd [49](#)

M

MAXSESSIONS server option [44](#)
MDisk [15](#)
mkdir command [30](#)
mkfs command [30](#)
mount command [30](#)
multipath I/O for disk storage [28](#)

N

NUMOPENVOLSALLOWED server option [44](#)

O

Operations Center [53, 79](#)

P

passwd command [49](#)
passwords
 default [49](#)
 updating [49](#)
performance
 evaluating [57](#)
 extra small system [57](#)
 large system [60](#)
 medium system [59](#)
 small system [58](#)
 testing [34](#)
 workload simulation tool [61](#)
performance results [57](#)
planning worksheet [10](#)

Q

QUERY EVENT server command [49](#)

R

RAID arrays [15](#)
Red Hat Enterprise Linux x86-64 [10](#)
REGISTER NODE server command [49](#)
registering nodes to the server [49](#)
replication storage rules [53](#)
replication subrules rules [53](#)

S

schedules
 client [41](#)
 predefined client [73](#)
 server [41](#)
script
 blueprint configuration script [41](#)
 configuration clean up [48](#)
 storage preparation [30](#)
 workload simulation tool [34](#)
server
 determining the size of [3](#)
 installing [38, 39](#)
 obtaining installation packages [38](#)
server commands
 DEFINE ASSOCIATION [49](#)
 QUERY EVENT [49](#)
 REGISTER NODE [49](#)
 SET ACTLOGRETENTION [43](#)
 SET EVENTRETENTION [43](#)
 SET MAXSCHEDSESSIONS [43](#)
server options
 ACTIVELOGSIZE [44](#)
 ALLOWREORGINDEX [45](#)
 ALLOWREORGTABLE [45](#)
 ARCHLOGCOMPRESS [44](#)
 COMMTIMEOUT [44](#)
 DEDUPDELETIONTHREADS [44](#)
 DEDUPREQUIRESBACKUP [44](#)
 DEVCONFIG [44](#)
 DIRECTIO [44](#)
 DISABLEREORGINDEX [45](#)
 DISABLEREORGTABLE [45](#)
 EXPINTERVAL [44](#)
 IDLETIMEOUT [44](#)
 MAXSESSIONS [44](#)
 NUMOPENVOLSALLOWED [44](#)
 REORGBEGINTIME [45](#)
 REORGDURATION [45](#)
 VOLUMEHISTORY [44](#)
 SET ACTLOGRETENTION server command [43](#)
 SET EVENTRETENTION server command [43](#)
 SET MAXSCHEDSESSIONS server command [43](#)
software prerequisites [10](#)
storage configuration
 planning for [10](#)
storage layout
 large system [20, 21, 31](#)
 medium system [18](#)
 small system [17](#)
storage preparation [29](#)
Storwize V5010 systems
 hardware requirements [7](#)
Storwize V5030 systems
 hardware requirements [8](#)
Supermicro SuperServer 2029U-E1CRT
 hardware requirements [7](#)
system setup [23](#)

T

tasks for configuration [23](#)
testing system performance [34](#)

total managed data [3](#)
troubleshooting [79](#)

U

ulimits [43](#)

V

VDisk [31](#)
virtual hardware requirements [6](#)
VOLUMEHISTORY server option [44](#)

W

What's new [vii](#)
workload simulation tool [34](#), [61](#)

