IBM Spectrum Protect
for Linux
Version 8.1.2

Installation Guide

IBM Spectrum Protect
for Linux
Version 8.1.2

Installation Guide

Note:
Before you use this information and the product it supports, read the information in "Notices" on page 171.

This edition applies to version 8, release 1, modification 2 of IBM Spectrum Protect (product numbers 5725-W98, 5725-W99, 5725-X15), and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication

This publication contains installation and configuration instructions for the IBM Spectrum Protect™ server, server languages, license, and device driver.

Instructions for installing the Operations Center are also included in this publication.

Who should read this guide

This publication is intended for system administrators who install, configure, or upgrade the IBM Spectrum Protect server or Operations Center.

Installable components

The IBM Spectrum Protect server and licenses are required components.

Table 1 describes all the installable components. These components are in several different installation packages.

Table 1. IBM Spectrum Protect installable components

IBM Spectrum Protect component	Description	Additional information
Server (required)	Includes the database, the Global Security Kit (GSKit), IBM® Java™ Runtime Environment (JRE), and tools to help you configure and manage the server.	See Chapter 2, "Installing the server components," on page 53.
Language package (optional)	Each language package (one for each language) contains language-specific information for the server.	See "Installing server language packages" on page 57.
Licenses (required)	Includes support for all licensed features. After you install this package, you must register the licenses you purchased.	Use the **REGISTER LICENSE** command.
Devices (optional)	Extends media management capability.	A list of devices that are supported by this driver is available from the IBM Support Portal.
Storage agent (optional)	Installs the component that allows client systems to write data directly to, or read data directly from, storage devices that are attached to a storage area network (SAN). **Remember:** IBM Spectrum Protect for Storage Area Networks is a separately licensed product.	For more information about storage agents, see Tivoli Storage Manager for Storage Area Networks (V7.1.1).

Table 1. IBM Spectrum Protect installable components (continued)

IBM Spectrum Protect component	Description	Additional information
Operations Center (optional)	Installs the Operations Center, which is a web-based interface for managing your storage environment.	See Part 2, "Installing and upgrading the Operations Center," on page 107.

Publications

The IBM Spectrum Protect product family includes IBM Spectrum Protect Snapshot, IBM Spectrum Protect for Space Management, IBM Spectrum Protect for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Knowledge Center.

What's new in this release

This release of IBM Spectrum Protect introduces new features and updates.

For a list of new features and updates, see What's new.

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

Part 1. Installing and upgrading the server

Install and upgrade the IBM Spectrum Protect server.

Chapter 1. Planning to install the server

Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to IBM Spectrum Protect server-managed storage.

What you should know first

Before installing IBM Spectrum Protect, be familiar with your operating systems, storage devices, communication protocols, and system configurations.

Server maintenance releases, client software, and publications are available from the IBM Support Portal.

Restriction: You can install and run the Version 8.1.2 server on a system that already has DB2® installed on it, whether DB2 was installed independently or as part of some other application, with some restrictions. For details, see "Compatibility of the IBM Spectrum Protect server with other DB2 products on the system" on page 28.

Experienced DB2 administrators can choose to perform advanced SQL queries and use DB2 tools to monitor the database. Do not, however, use DB2 tools to change DB2 configuration settings from those that are preset by IBM Spectrum Protect, or alter the DB2 environment for IBM Spectrum Protect in other ways, such as with other products. The V8.1.2 server has been built and tested extensively using the data definition language (DDL) and database configuration that the server deploys.

Attention: Do not alter the DB2 software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

What you should know about security before you install or upgrade the server

Before you install IBM Spectrum Protect V8.1.2 or later, review information about the enhanced security features and the requirements for updating your environment.

About this task

Security enhancements that were introduced in V8.12 and later enforce stricter security settings. To ensure that communication between servers and clients is not interrupted when you install or upgrade IBM Spectrum Protect software to V8.1.2, follow the procedure.

Procedure

1. Install or upgrade the IBM Spectrum Protect servers to 8.1.2 or later.
2. Install or upgrade the backup-archive clients. For more information, see Installing and configuring clients.

   For information about scheduling deployment of client updates from the server, see the following documents:

- For IBM Spectrum Protect 8.1.2 or later servers, see technote 2004596.
- For IBM Tivoli® Storage Manager V7.1 servers and IBM Spectrum Protect V8.1.0 and V8.1.1 servers, see technote 1673299.

3. Configure the options for backup-archive clients. For more information, see Upgrading the IBM Spectrum Protect Server and the IBM Spectrum Protect Client.

Planning for optimal performance

Before you install the IBM Spectrum Protect server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.

Procedure

1. Review "What you should know first" on page 3.
2. Review each of the following sub-sections.

Planning for the server hardware and the operating system

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
Does the operating system and hardware meet or exceed requirements? • Number and speed of processors • System memory • Supported operating system level	If you are using the minimum required amount of memory, you can support a minimal workload. You can experiment by adding more system memory to determine whether the performance is improved. Then, decide whether you want to keep the system memory dedicated to the server. Test the memory variations by using the entire daily cycle of the server workload. If you run multiple servers on the system, add the requirements for each server to get the requirements for the system.	Review operating system requirements at technote 1243309. Additionally, review the guidance in Tuning tasks for operating systems and other applications. For more information about requirements when these features are in use, see the following topics: • Checklist for data deduplication • Checklist for node replication For more information about sizing requirements for the server and storage, see the IBM Spectrum Protect Blueprint.
Are disks configured for optimal performance?	The amount of tuning that can be done for different disk systems varies. Ensure that the appropriate queue depths and other disk system options are set.	For more information, see the following topics: • "Planning for server database disks" • "Planning for server recovery log disks" • "Planning for storage pools in DISK or FILE device classes"

Question	Tasks, characteristics, options, or settings	More information
Does the server have enough memory?	Heavier workloads and advanced features such as data deduplication and node replication require more than the minimum system memory that is specified in the system requirements document. For databases that are not enabled for data deduplication, use the following guidelines to specify memory requirements: • For databases less than 500 GB, you need 16 GB of memory. • For databases with a size of 500 GB - 1 TB, you need 24 GB of memory. • For databases with a size of 1 TB - 1.5 TB, you need 32 GB of memory. • For databases greater than 1.5 TB, you need 40 GB of memory. Ensure that you allocate extra space for the active log and the archive log for replication processing.	For more information about requirements when these features are in use, see the following topics: • Checklist for data deduplication • Checklist for node replication • Memory requirements
Does the system have enough host bus adapters (HBAs) to handle the data operations that the IBM Spectrum Protect server must run simultaneously?	Understand what operations require use of HBAs at the same time. For example, a server must store 1 GB/sec of backup data while also doing storage pool migration that requires 0.5 GB/sec capacity to complete. The HBAs must be able to handle all of the data at the speed required.	See Tuning HBA capacity.

Question	Tasks, characteristics, options, or settings	More information
Is network bandwidth greater than the planned maximum throughput for backups?	Network bandwidth must allow the system to complete operations such as backups in the time that is allowed or that meets service level commitments. For node replication, network bandwidth must be greater than the planned maximum throughput.	For more information, see the following topics: • Tuning network performance • Checklist for node replication
Are you using a preferred file system for IBM Spectrum Protect server files?	Use a file system that ensures optimal performance and data availability. The server uses direct I/O with file systems that support the feature. Using direct I/O can improve throughput and reduce processor use. The following list identifies the preferred file system: • Use either the ext3, ext4, or xfs file system for the database, recovery log, and storage pool data. Use the following file system that is appropriate for your operating system and level: – For Red Hat Enterprise Linux x86_64, use the ext3, ext4, or xfs file system. If Red Hat Enterprise Linux 6.4 or later is installed, use the ext4 or xfs file system. – For SUSE Linux Enterprise Server and for Red Hat Enterprise Linux ppc64, use the ext3 or xfs file system. Using xfs on SUSE Linux Enterprise Server 12 requires kernel-default-3.12.32-33.1.x86_64.rpm or later.	For more information, see Configuring the operating system for disk performance.

Question	Tasks, characteristics, options, or settings	More information
Are you planning to configure enough paging space?	Paging space, or swap space, extends the memory that is available for processing. When the amount of free RAM in the system is low, programs or data that is not in use are moved from memory to paging space. This action releases memory for other activities, such as database operations. Use a minimum of 32 GB of paging space or 50% of your RAM, whichever value is larger.	
Are you planning to tune the kernel parameters after installation of the server?	You must tune kernel parameters.	See the information about tuning kernel parameters: Linux: Tuning kernel parameters for Linux systems

Planning for the server database disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
Is the database on fast, low-latency disks?	Do not use the following drives for the IBM Spectrum Protect database: • Nearline SAS (NL-SAS) • Serial Advanced Technology Attachment (SATA) • Parallel Advanced Technology Attachment (PATA) Do not use internal disks that are included by default in most server hardware. Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance. If you plan to use the data deduplication functions of IBM Spectrum Protect, focus on disk performance in terms of I/O operations per second (IOPS).	For more information, see Checklist for data deduplication.

Question	Tasks, characteristics, options, or settings	More information
Is the database stored on disks or LUNs that are separate from disks or LUNs that are used for the active log, archive log, and storage pool volumes?	Separation of the server database from other server components helps reduce contention for the same resources by different operations that must run at the same time. **Tip:** The database and the archive log can share an array when you use solid-state drive (SSD) technology.	
If you are using RAID, do you know how to select the optimal RAID level for your system? Are you defining all LUNs with the same size and type of RAID?	When a system must do large numbers of writes, RAID 10 outperforms RAID 5. However, RAID 10 requires more disks than RAID 5 for the same amount of usable storage. If your disk system is RAID, define all your LUNs with the same size and type of RAID. For example, do not mix 4+1 RAID 5 with 4+2 RAID 6.	
If an option to set the strip size or segment size is available, are you planning to optimize the size when you configure the disk system?	If you can set the strip size or segment size, use 64 KB or 128 KB sizes on disk systems for the database.	The block size that is used for the database varies depending on the table space. Most table spaces use 8 KB blocks, but some use 32 KB blocks.

Question	**Tasks, characteristics, options, or settings**	**More information**
Are you planning to create at least four directories, also called storage paths, on four separate LUNs for the database? Create one directory per distinct array on the subsystem. If you have fewer than three arrays, create a separate LUN volume within the array.	Heavier workloads and use of some features require more database storage paths than the minimum requirements. Server operations such as data deduplication drive a high number of input/output operations per second (IOPS) for the database. Such operations perform better when the database has more directories. For server databases that are larger than 2 TB or are expected to grow to that size, use eight directories. Consider planned growth of the system when you determine how many storage paths to create. The server uses the higher number of storage paths more effectively if the storage paths are present when the server is first created. Use the *DB2_PARALLEL_IO* variable to force parallel I/O to occur on table spaces that have one container, or on table spaces that have containers on more than one physical disk. If you do not set the *DB2_PARALLEL_IO* variable, I/O parallelism is equal to the number of containers that are used by the table space. For example, if a table space spans four containers, the level of I/O parallelism that is used is 4.	For more information, see the following topics: • Checklist for data deduplication • Checklist for node replication For help with forecasting growth when the server deduplicates data, see technote 1596944. For the most recent information about database size, database reorganization, and performance considerations for IBM Spectrum Protect servers, see technote 1683633. For information about setting the *DB2_PARALLEL_IO* variable, see Recommended settings for IBM DB2 registry variables.
Are all directories for the database the same size?	Directories that are all the same size ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching. This guideline also applies if you must add storage paths after the initial configuration of the server.	
Are you planning to raise the queue depth of the database LUNs on AIX® systems?	The default queue depth is often too low.	See Configuring AIX systems for disk performance.

Planning for the server recovery log disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
Are the active log and archive log stored on disks or LUNs that are separate from what is used for the database and storage pool volumes?	Ensure that the disks where you place the active log are not used for other server or system purposes. Do not place the active log on disks that contain the server database, the archive log, or system files such as page or swap space.	Separation of the server database, active log, and archive log helps to reduce contention for the same resources by different operations that must run at the same time.
Are the logs on disks that have nonvolatile write cache?	Nonvolatile write cache allows data to be written to the logs as fast as possible. Faster write operations for the logs can improve performance for server operations.	
Are you setting the logs to a size that adequately supports the workload?	If you are not sure about the workload, use the largest size that you can. **Active log** The maximum size is 512 GB, set with the **`ACTIVELOGSIZE`** server option. Ensure that there is at least 8 GB of free space on the active log file system after the fixed size active logs are created. **Archive log** The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Make the archive log at least as large as the active log.	• For log sizing details, see the recovery log information in technote 1421060. • For information about sizing when you use data deduplication, see Checklist for data deduplication.
Are you defining an archive failover log? Are you placing this log on a disk that is separate from the archive log?	The archive failover log is for emergency use by the server when the archive log becomes full. Slower disks can be used for the archive failover log.	Use the **`ARCHFAILOVERLOGDIRECTORY`** server option to specify the location of the archive failover log. Monitor the usage of the directory for the archive failover log. If the archive failover log must be used by the server, the space for the archive log might not be large enough.

Question	Tasks, characteristics, options, or settings	More information
If you are mirroring the active log, are you using only one type of mirroring?	You can mirror the log by using one of the following methods. Use only one type of mirroring for the log. • Use the **`MIRRORLOGDIRECTORY`** option that is available for the IBM Spectrum Protect server to specify a mirror location. • Use software mirroring, such as Logical Volume Manager (LVM) on AIX. • Use mirroring in the disk system hardware.	If you mirror the active log, ensure that the disks for both the active log and the mirror copy have equal speed and reliability. For more information, see Configuring and tuning the recovery log.

Planning for directory-container and cloud-container storage pools

Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.

Question	Tasks, characteristics, options, or settings	More information
Measured in terms of input/output operations per second (IOPS), are you using fast disk storage for the IBM Spectrum Protect database?	Use a high-performance disk for the database. Use solid-state drive technology for data deduplication processing. Ensure that the database has a minimum capability of 3000 IOPS. For each TB of data that is backed up daily (before data deduplication), add 1000 IOPS to this minimum. For example, an IBM Spectrum Protect server that is ingesting 3 TB of data per day would need 6000 IOPS for the database disks: `3000 IOPS minimum + 3000 (3 TB x 1000 IOPS) = 6000 IOPS`	For recommendations about disk selection, see "Planning for server database disks". For more information about IOPS, see the IBM Spectrum Protect Blueprints.

Question	Tasks, characteristics, options, or settings	More information
Do you have enough memory for the size of your database?	Use a minimum of 40 GB of system memory for IBM Spectrum Protect servers, with a database size of 100 GB, that are deduplicating data. If the retained capacity of backup data grows, the memory requirement might need to be higher. Monitor memory usage regularly to determine whether more memory is required. Use more system memory to improve caching of database pages. The following memory size guidelines are based on the daily amount of new data that you back up: • 128 GB of system memory for daily backups of data, where the database size is 1 - 2 TB • 192 GB of system memory for daily backups of data, where the database size is 2 - 4 TB	Memory requirements
Have you properly sized the storage capacity for the database active log and archive log?	Configure the server to have a minimum active log size of 128 GB by setting the **ACTIVELOGSIZE** server option to a value of 131072. The suggested starting size for the archive log is 1 TB. The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Ensure that there is at least 10% extra disk space for the file system than the size of the archive log. Use a directory for the database archive logs with an initial free capacity of at least 1 TB. Specify the directory by using the **ARCHLOGDIRECTORY** server option. Define space for the archive failover log by using the **ARCHFAILOVERLOGDIRECTORY** server option.	For more information about sizing for your system, see the IBM Spectrum Protect Blueprints.

Question	Tasks, characteristics, options, or settings	More information
Is compression enabled for the archive log and database backups?	Enable the `ARCHLOGCOMPRESS` server option to save storage space. This compression option is different from inline compression. Inline compression is enabled by default with IBM Spectrum Protect V7.1.5 and later. **Restriction:** Do not use this option if the amount of backed up data exceeds 6 TB per day.	For more information about compression for your system, see the IBM Spectrum Protect Blueprints.
Are the IBM Spectrum Protect database and logs on separate disk volumes (LUNs)? Is the disk that is used for the database configured according to best practices for a transactional database?	The database must not share disk volumes with IBM Spectrum Protect database logs or storage pools, or with any other application or file system.	For more information about server database and recovery log configuration, see Server database and recovery log configuration and tuning.
Are you using a minimum of eight (2.2 GHz or equivalent) processor cores for each IBM Spectrum Protect server that you plan to use with data deduplication?	If you are planning to use client-side data deduplication, verify that client systems have adequate resources available during a backup operation to complete data deduplication processing. Use a processor that is at least the minimum equivalent of one 2.2 GHz processor core per backup process with client-side data deduplication.	• Effective planning and use of deduplication • IBM Spectrum Protect Blueprints
Did you allocate enough storage space for the database?	For a rough estimate, plan for 100 GB of database storage for every 50 TB of data that is to be protected in deduplicated storage pools. *Protected data* is the amount of data before data deduplication, including all versions of objects stored. As a best practice, define a new container storage pool exclusively for data deduplication. Data deduplication occurs at the storage-pool level, and all data within a storage pool, except encrypted data, is deduplicated.	

Question	Tasks, characteristics, options, or settings	More information
Have you estimated storage pool capacity to configure enough space for the size of your environment?	You can estimate capacity requirements for a deduplicated storage pool by using the following technique: 1. Estimate the base size of the source data. 2. Estimate the daily backup size by using an estimated change and growth rate. 3. Determine retention requirements. 4. Estimate the total amount of source data by factoring in the base size, daily backup size, and retention requirements. 5. Apply the deduplication ratio factor. 6. Apply the compression ratio factor. 7. Round up the estimate to consider transient storage pool usage.	For an example of using this technique, see Effective planning and use of deduplication.
Have you distributed disk I/O over many disk devices and controllers?	Use arrays that consist of as many disks as possible, which is sometimes referred to as wide striping. Ensure that you use one database directory per distinct array on the subsystem. Set the *DB2_PARALLEL_IO* registry variable to enable parallel I/O for each table space used if the containers in the table space span multiple physical disks. When I/O bandwidth is available and the files are large, for example 1 MB, the process of finding duplicates can occupy the resources of an entire processor. When files are smaller, other bottlenecks can occur. Specify eight or more file systems for the deduplicated storage pool device class so that I/O is distributed across as many LUNs and physical devices as possible.	For guidelines about setting up storage pools, see "Planning for storage pools in DISK or FILE device classes". For information about setting the *DB2_PARALLEL_IO* variable, see Recommended settings for IBM DB2 registry variables.
Have you scheduled daily operations based on your backup strategy?	The best practice sequence of operations is in the following order: 1. Client backup 2. Storage pool protection 3. Node replication 4. Database backup 5. Expire inventory	• Scheduling data deduplication and node replication processes • Daily operations for directory-container storage pools

Question	Tasks, characteristics, options, or settings	More information
Do you have enough storage to manage the DB2 lock list?	If you deduplicate data that includes large files or large numbers of files concurrently, the process can result in insufficient storage space. When the lock list storage is insufficient, backup failures, data management process failures, or server outages can occur. File sizes greater than 500 GB that are processed by data deduplication are most likely to deplete storage space. However, if many backup operations use client-side data deduplication, this problem can also occur with smaller-sized files.	For information about tuning the DB2 **LOCKLIST** parameter, see Tuning server-side data deduplication.
Is sufficient bandwidth available to transfer data to an IBM Spectrum Protect server?	To transfer data to an IBM Spectrum Protect server, use client-side or server-side data deduplication and compression to reduce the bandwidth that is required. Use a V7.1.5 server or higher to use inline compression and use a V7.1.6 or later client to enable enhanced compression processing.	For more information, see the **enablededup** client option.
Have you determined how many storage pool directories to assign to each storage pool?	Assign directories to a storage pool by using the **DEFINE STGPOOLDIRECTORY** command. Create multiple storage pool directories and ensure that each directory is backed up to a separate disk volume (LUN).	
Did you allocate enough disk space in the cloud-container storage pool?	To prevent backup failures, ensure that the local directory has enough space. Use the following list as a guide for optimal disk space: • For serial-attached SCSI (SAS) and spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount, in terabytes, for disk space. • Provide 3 TB for flash-based storage systems with fast network connections to on-premises, high-performance cloud systems. • Provide 5 TB for solid-state drive (SSD) systems with fast network connections to high-performance cloud systems.	

Question	Tasks, characteristics, options, or settings	More information
Did you select the appropriate type of local storage?	Ensure that data transfers from local storage to cloud finish before the next backup cycle starts. **Tip:** Data is removed from local storage soon after it moves to the cloud. Use the following guidelines: • Use flash or SSD for large systems that have high-performing cloud systems. Ensure that you have a dedicated 10 GB wide area network (WAN) link with a high-speed connection to the object storage. For example, use flash or SSD if you have a dedicated 10 GB WAN link plus a high-speed connection to either an IBM Cloud Object Storage location or to an Amazon Simple Storage Service (Amazon S3) data center. • Use larger capacity 15000 rpm SAS disks for these scenarios: – Medium-sized systems – Slower cloud connections, for example, 1 GB – When you use IBM Cloud Object Storage as your service provider across several regions • For SAS or spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount for disk space, in terabytes.	

Planning for storage pools in DISK or FILE device classes

Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use `DISK` or `FILE` device classes.

Question	Tasks, characteristics, options, or settings	More information
Can the storage pool LUNs sustain throughput rates for 256 KB sequential reads and writes to adequately handle the workload within the time constraints?	When you are planning for peak loads, consider all the data that you want the server to read or write to the disk storage pools simultaneously. For example, consider the peak flow of data from client backup operations and server data-movement operations such as migration that run at the same time. The IBM Spectrum Protect server reads and writes to storage pools predominantly in 256 KB blocks. If the disk system includes the capability, configure the disk system for optimal performance with sequential read/write operations rather than random read/write operations.	For more information, see Analyzing the basic performance of disk systems.
Is the disk configured to use read and write cache?	Use more cache for better performance.	
For storage pools that use `FILE` device classes, have you determined a good size to use for the storage pool volumes?	Review the information in Optimal number and size of volumes for storage pools that use disk. If you do not have the information to estimate a size for `FILE` device class volumes, start with volumes that are 50 GB.	Typically, problems arise more frequently when the volumes are too small. Few problems are reported when volumes are larger than needed. When you determine the volume size to use, as a precaution choose a size that might be larger than necessary.
For storage pools that use `FILE` device classes, are you using preallocated volumes?	Scratch volumes can cause file fragmentation. To ensure that a storage pool does not run out of volumes, set the **MAXSCRATCH** parameter to a value greater than zero.	Use the **DEFINE VOLUME** server command to preallocate volumes in the storage pool. Use the **DEFINE STGPOOL** or **UPDATE STGPOOL** server command to set the **MAXSCRATCH** parameter.
For storage pools that use `FILE` device classes, have you compared the maximum number of client sessions to the number of volumes that are defined?	Always maintain enough usable volumes in the storage pools to allow for the expected peak number of client sessions that run at one time. The volumes might be scratch volumes, empty volumes, or partly filled volumes.	For storage pools that use `FILE` device classes, only one session or process can write to a volume at the same time.

Question	Tasks, characteristics, options, or settings	More information
For storage pools that use `FILE` device classes, have you set the **`MOUNTLIMIT`** parameter of the device class to a value that is high enough to account for the number of volumes that might be mounted in parallel?	For storage pools that use data deduplication, the **`MOUNTLIMIT`** parameter is typically in the range of 500 - 1000. Set the value for **`MOUNTLIMIT`** to the maximum number of mount points that are needed for all active sessions. Consider parameters that affect the maximum number of mount points that are needed: • The **`MAXSESSIONS`** server option, which is the maximum number of IBM Spectrum Protect sessions that can run concurrently. • The **`MAXNUMMP`** parameter, which sets the maximum number of mount points that each client node can use. For example, if the maximum number of client node backup sessions is typically 100 and each of the nodes has `MAXNUMMP=2`, multiply 100 nodes by the 2 mount points for each node to get the value of 200 for the **`MOUNTLIMIT`** parameter.	Use the **`REGISTER NODE`** or **`UPDATE NODE`** server command to set the **`MAXNUMMP`** parameter for client nodes.
For storage pools that use `DISK` device classes, have you determined how many storage pool volumes to put on each file system?	How you configure the storage for a storage pool that uses a `DISK` device class depends on whether you are using RAID for the disk system. If you are not using RAID, then configure one file system per physical disk, and define one storage pool volume for each file system. If you are using RAID 5 with *n* + 1 volumes, configure the storage in one of the following ways: • Configure *n* file systems on the LUN and define one storage pool volume per file system. • Configure one file system and *n* storage pool volumes for the LUN.	For an example layout that follows this guideline, see Sample layout of server storage pools.
Did you create your storage pools to distribute I/O across multiple file systems?	Ensure that each file system is on a different LUN on the disk system. Typically, having 10 - 30 file systems is a good goal, but ensure that the file systems are no smaller than approximately 250 GB.	For details, see the following topics: • Tuning disk storage for the server • Tuning and configuring storage pools and volumes

Planning for the correct type of storage technology

Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with IBM Spectrum Protect.

Procedure

Review the following table to help you to choose the correct type of storage technology for the storage resources that the server requires.

Table 2. Storage technology types for IBM Spectrum Protect storage requirements

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
Solid-state disk (SSD)	Place the database on SSD in the following circumstances: • You are using IBM Spectrum Protect data deduplication. • You are backing up more than 8 TB of new data daily.	If you place the IBM Spectrum Protect database on an SSD, as a best practice, place the active log on an SSD. If space is not available, use high-performance disk instead.	Save SSDs for use with the database and active log. The archive log and archive failover logs can be placed on slower storage technology types.	Save SSDs for use with the database and active log. Storage pools can be placed on slower storage technology types.
High-performance disk with the following characteristics: • 15k rpm disk • Fibre Channel or serial-attached SCSI (SAS) interface	Use high-performance disks in the following circumstances: • The server does not do data deduplication. • The server does not do node replication. Isolate the server database from its logs and storage pools, and from data for other applications.	Use high-performance disks in the following circumstances: • The server does not do data deduplication. • The server does not do node replication. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use high-performance disks for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use high-performance disks for storage pools in the following circumstances: • Data is frequently read. • Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.
Medium-performance or high-performance disk with the following characteristics: • 10k rpm disk • Fibre Channel or SAS interface	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. Isolate the server database from its logs and storage pools, and from data for other applications.	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use medium-performance or high-performance disk for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use medium-performance or high-performance disk for storage pools in the following circumstances: • Data is frequently read. • Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.

Table 2. Storage technology types for IBM Spectrum Protect storage requirements (continued)

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
SATA, network-attached storage	Do not use this storage for the database. Do not place the database on XIV storage systems.	Do not use this storage for the active log.	Use of this slower storage technology is acceptable because these logs are written once and infrequently read.	Use this slower storage technology in the following circumstances: • Data is infrequently written, for example written once. • Data is infrequently read. .
Tape and virtual tape				Use for long-term retention or if data is infrequently used.

Applying best practices to the server installation

Typically, hardware configuration and selection have the most significant effect on the performance of an IBM Spectrum Protect solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Spectrum Protect.

Procedure

- The following best practices are the most important for optimal performance and problem prevention.
- Review the table to determine the best practices that apply to your environment.

Best practice	More information
Use fast disks for the server database. Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.	Use fast, low-latency disks for the database. Using SSD is essential if you are using data deduplication and node replication. Avoid Serial Advanced Technology Attachment (SATA) and Parallel Advanced Technology Attachment (PATA) disks. For details and more tips, see the following topics: • "Planning for server database disks" • "Planning for the correct type of storage technology"
Ensure that the server system has enough memory.	Review operating system requirements in technote 1243309. Heavier workloads require more than the minimum requirements. Advanced features such as data deduplication and node replication can require more than the minimum memory that is specified in the system requirements document. If you plan to run multiple instances, each instance requires the memory that is listed for one server. Multiply the memory for one server by the number of instances that are planned for the system.

Best practice	More information
Separate the server database, the active log, the archive log, and disk storage pools from each other.	Keep all IBM Spectrum Protect storage resources on separate disks. Keep storage pool disks separate from the disks for the server database and logs. Storage pool operations can interfere with database operations when both are on the same disks. Ideally, the server database and logs are also separated from each other. For details and more tips, see the following topics: • "Planning for server database disks" • "Planning for server recovery log disks" • "Planning for storage pools in DISK or FILE device classes"
Use at least four directories for the server database. For larger servers or servers that use advanced features, use eight directories.	Place each directory on a LUN that is isolated from other LUNs and from other applications. A server is considered to be large if its database is larger than 2 TB or is expected to grow to that size. Use eight directories for such servers. See "Planning for server database disks".
If you are using data deduplication, node replication, or both, follow the guidelines for database configuration and other items.	Configure the server database according to the guidelines, because the database is extremely important to how well the server runs when these features are being used. For details and more tips, see the following topics: • Checklist for data deduplication • Checklist for node replication
For storage pools that use `FILE` type device classes, follow the guidelines for the size of storage pool volumes. Typically, 50 GB volumes are best.	Review the information in Optimal number and size of volumes for storage pools that use disk to help you to determine volume size. Configure storage pool devices and file systems based on throughput requirements, not only on capacity requirements. Isolate the storage devices that are used by IBM Spectrum Protect from other applications that have high I/O, and ensure that there is enough throughput to that storage. For more details, see Checklist for storage pools on DISK or FILE.
Schedule IBM Spectrum Protect client operations and server maintenance activities to avoid or minimize overlap of operations.	For more details, see the following topics: • Tuning the schedule for daily operations • Checklist for server configuration
Monitor operations constantly.	By monitoring, you can find problems early and more easily identify causes. Keep records of monitoring reports for up to a year to help you identify trends and plan for growth. See Monitoring and maintaining the environment for performance.

Minimum system requirements

To install the IBM Spectrum Protect server on a Linux system, it is necessary to have a minimum level of hardware and software, including a communication method and the most current device driver.

These tables list the minimum hardware and software requirements for the installation of an IBM Spectrum Protect server. Use these requirements as a starting point for systems without data deduplication. The optimal IBM Spectrum Protect environment is set up with data deduplication by using the IBM Spectrum Protect Blueprints. For the most current information about system requirements, see technote 1243309.

The IBM Spectrum Protect device driver package does not contain a device driver for this operating system because a SCSI generic device driver is used. Configure the device driver before using the IBM Spectrum Protect server with tape devices. The IBM Spectrum Protect driver package contains driver tools and ACSLS daemons. You can locate IBM driver packages at the Fix Central website.

Requirements, supported devices, client installation packages, and fixes are available in the IBM Support Portal for IBM Spectrum Protect. After you install IBM Spectrum Protect and before you customize it for your use, go to the website and download and apply any applicable fixes.

Minimum Linux X86_64 server requirements

Before you install an IBM Spectrum Protect server on a Linux X86_64 operating system, review the hardware and software requirements.

Hardware requirements

Table 3 describes the minimum hardware requirements for the server. If the server does not meet the minimum requirements, the installation fails. For more details about planning disk space, see “Capacity planning” on page 31.

Table 3. Hardware requirements

Type of hardware	Hardware requirements
Server	An AMD64 or Intel EMT-64 processor
Disk space	The following minimum values for disk space: • 5 GB for the installation directory • 512 MB for the /var directory • 2 GB for the `/tmp` directory • 128 MB in the home directory for the root user. • 2 GB for the shared resources area In case a problem arises and any diagnosis is needed, it is optimal to have temporary or other space available on the system for a first failure data capture (FFDC) log or for other temporary uses such as for collecting trace logs. Significant additional disk space is required for database and log files. The size of the database depends on the number of client files to be stored and the method by which the server manages them. The default active log space is 16 GB, the minimum that is needed for most workloads and configurations. When you create the active log, you need at least 64 GB to run replication. If replication and data deduplication are both being used, create an active log of 128 GB. Allocate at least three times the default active log space for the archive log (48 GB). Ensure that you have sufficient resources if you are using data deduplication or expect a heavy client workload. For optimal performance and to facilitate I/O, specify at least two equally sized containers or Logical Unit Numbers (LUNs) for the database. In addition, each active log and archive log needs its own container or LUN. Ensure that you see “Capacity planning” on page 31 for more details about disk space.

Table 3. Hardware requirements (continued)

Type of hardware	Hardware requirements
Memory	The following minimum values for memory: • 16 GB for standard server operations without data deduplication and node replication • 24 GB for data deduplication or node replication • 32 GB for node replication with data deduplication For more specific memory requirements for larger databases and higher ingestion capability, see the IBM Spectrum Protect server memory tuning table. For more specific memory requirements when you are using data deduplication, see the IBM Spectrum Protect Blueprint for your operating system.

Software requirements

Table 4 describes the minimum software requirements that are needed for a server on a Linux X86_64 system.

Table 4. Software requirements

Type of software	Minimum software requirements
Operating system	The IBM Spectrum Protect server on Linux X86_64 requires one of the following operating systems: • Red Hat Enterprise Linux 6.7 • Red Hat Enterprise Linux 7, including updates • SUSE Linux Enterprise Server 11, Service Pack 4 or later • SUSE Linux Enterprise Server 12
Libraries	GNU C libraries, Version 2.3.3-98.38 or later, which are installed on the IBM Spectrum Protect system. For SUSE Linux Enterprise Servers: • libaio • libstdc++.so.6 at version 4.3 or later (32-bit and 64-bit packages are required) For Red Hat Enterprise Linux Servers: • libaio • libstdc++.so.6 (32 and 64 bit packages are required) • numactl.x86_64 To determine if SELinux is installed and in enforcing mode, take 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. To disable SELinux, complete one of the following tasks: • Set permissive mode by issuing the `setenforce 0` command as a superuser. • Modify the `/etc/sysconfig/selinux` file and restart the machine.

Table 4. Software requirements (continued)

Type of software	Minimum software requirements
Communication protocol	• TCP/IP Version 4 or Version 6, which is standard with Linux • Shared memory protocol (with IBM Spectrum Protect Linux X86_64 client)
Processing	Asynchronous I/O must be enabled. On Linux kernels at 2.6 or later, install the libaio library to enable asynchronous I/O.
Device drivers	The IBM Spectrum Protect passthru device driver is used for non-IBM devices. It uses the SCSI passthru interface to communicate with tape devices and tape libraries. The Linux SCSI Generic (sg) device driver is required for tape drives and tape libraries. The IBM Spectrum Protect device driver package contains device driver tools and ACSLS daemons. For the IBM 3590, 3592, or the Ultrium tape library or drives, the IBM device drivers are required. Install the most current device drivers. You can locate IBM driver packages at Fix Central. Configure the device drivers before you use the server with tape devices.
Other software	Korn Shell (ksh) is required. Configure the I/O completion ports (IOCP) on the operating system. To authenticate IBM Spectrum Protect users with a Lightweight Directory Access Protocol (LDAP) server, you must use one of the following directory servers: • Microsoft Active Directory (Windows Server 2012, Windows Server 2012 R2) • IBM Security Directory Server V6.3 • IBM Security Directory Server V6.4

Minimum Linux on System z server requirements

Before you install an IBM Spectrum Protect server on a Linux on System z® operating system, review the hardware and software requirements.

Hardware requirements

Table 5 describes the minimum hardware requirements that are needed for your IBM Spectrum Protect Linux on System z system. For more details about planning disk space, see "Capacity planning" on page 31.

Table 5. Hardware requirements

Type of hardware	Hardware requirements
Server	An IBM zSeries, IBM System z9®, IBM System z10®, or IBM zEnterprise® System (z114 and z196) 64-bit native logical partition (LPAR) or z/VM® guest.

Table 5. Hardware requirements (continued)

Type of hardware	Hardware requirements
Disk space	The following minimum values for disk space: • 5 GB for the installation directory • 512 MB for the `/var` directory • 2 GB for the `/tmp` directory • 128 MB in the home directory for the root user. • 2 GB for the shared resources area In case a problem arises and any diagnosis is needed, it is optimal to have temporary or other space available on the system for a first failure data capture (FFDC) log or for other temporary uses such as for collecting trace logs. Significant additional disk space is required for database and log files. The size of the database depends on the number of client files to be stored and the method by which the server manages them. The default active log space is 16 GB, the minimum that is needed for most workloads and configurations. When you create the active log, you need at least 64 GB to run replication. If replication and data deduplication are both being used, create an active log of 128 GB. Allocate at least three times the default active log space for the archive log (48 GB). Ensure that you have sufficient resources if you are using data deduplication or expect a heavy client workload. For optimal performance and to facilitate I/O, specify at least two equally sized containers or Logical Unit Numbers (LUNs) for the database. In addition, each active log and archive log needs its own container or LUN. Ensure that you see “Capacity planning” on page 31 for more details about disk space.
Memory	The following minimum values for memory: • 16 GB for standard server operations without data deduplication and node replication • 24 GB for data deduplication or node replication • 32 GB for node replication with data deduplication For more specific memory requirements for larger databases and higher ingestion capability, see the IBM Spectrum Protect server memory tuning table. For more specific memory requirements when you are using data deduplication, see the IBM Spectrum Protect Blueprint for your operating system.

Software requirements

Table 6 on page 26 describes the minimum software requirements that are needed for your IBM Spectrum Protect Linux on System z system.

Table 6. Software requirements

Type of software	Minimum software requirements
Server	The IBM Spectrum Protect server on Linux on System z (s390x 64-bit architecture) requires one of the following operating systems: • Red Hat Enterprise Linux 7.1 • SUSE Linux Enterprise Server 12
Libraries	A GNU C library, Version 2.4-31.43.6, is installed on the IBM Spectrum Protect system. For SUSE Linux Enterprise Servers: • libaio • libstdc++.so.6 at version 4.3 or later (32-bit and 64-bit packages are required) • libxlc-1.2.0.0.151119a.s390x or later For Red Hat Enterprise Linux Servers: • libaio • libstdc++.so.6 (32-bit and 64-bit packages are required) • numactl.x86_64 • libxlc-1.2.0.0.151119a.s390x or later
Communication protocol	• TCP/IP Version 4 or Version 6, which is standard with Linux • Shared memory protocol (with IBM Spectrum Protect Version 8.1.2 Linux on System z client)
Processing	Asynchronous I/O must be enabled. On Linux kernels at 2.6 or later, install the libaio library to enable asynchronous I/O.
Other software	Korn Shell (ksh) is required. Configure the I/O completion ports (IOCP) on the operating system. To authenticate IBM Spectrum Protect users with a Lightweight Directory Access Protocol (LDAP) server, you must use one of the following directory servers: • Microsoft Active Directory (Windows Server 2012, Windows Server 2012 R2) • IBM Security Directory Server V6.3 • IBM Security Directory Server V6.4

Minimum Linux on Power Systems™ (little endian) server requirements

Before you install an IBM Spectrum Protect server on a Linux on Power Systems (little endian) operating system, review the hardware and software requirements.

Hardware requirements

Table 7 on page 27 describes the minimum hardware requirements for the server. If the server does not meet the minimum requirements, the installation fails. For more details about planning disk space, see “Capacity planning” on page 31.

Table 7. Hardware requirements

Type of hardware	Hardware requirements
Server	A Linux on Power Systems (little endian) server on an IBM system, such as one listed on the Linux on IBM Power Systems website.
Disk space	The following minimum disk space: • 5 GB for the installation directory • 128 MB in the home directory for the root user. • 2 GB for the shared resources area In case a problem arises and any diagnosis is needed, it is optimal to have temporary or other space available on the system for a first failure data capture (FFDC) log or for other temporary uses such as for collecting trace logs. Significant additional disk space is required for database and log files. The size of the database depends on the number of client files to be stored and the method by which the server manages them. The default active log space is 16 GB, the minimum that is needed for most workloads and configurations. When you create the active log, you need at least 64 GB to run replication. If replication and data deduplication are both being used, create an active log of 128 GB. Allocate at least three times the default active log space for the archive log (48 GB). Ensure that you have sufficient resources if you are using data deduplication or expect a heavy client workload. For optimal performance and to facilitate I/O, specify at least two equally sized containers or Logical Unit Numbers (LUNs) for the database. In addition, each active log and archive log needs its own container or LUN. Ensure that you see “Capacity planning” on page 31 for more details about disk space.
Memory	• 16 GB for standard server operations without data deduplication and node replication • 24 GB for data deduplication or node replication • 32 GB for node replication with data deduplication For more specific memory requirements for larger databases and higher ingestion capability, see the IBM Spectrum Protect server memory tuning table. For more specific memory requirements when you are using data deduplication, see the IBM Spectrum Protect Blueprint for your operating system.

Software requirements

Table 8 describes the minimum software requirements that are needed for your system.

Table 8. Software requirements

Type of software	Minimum software requirements
Operating system	The Red Hat Enterprise Linux (RHEL) 7.3 operating system with the PPC64LE architecture.

Table 8. Software requirements (continued)

Type of software	Minimum software requirements
Libraries	GNU C libraries, Version 2.4-31.30 and later. libaio.so.1 (32-bit and 64-bit packages).
Communication protocol	• TCP/IP Version 4 or Version 6, which is standard with Linux • Shared memory protocol (with a Version 8.1.2 client)
Processing	Asynchronous I/O must be enabled. On Linux kernels at 2.6 or later, install the libaio library to enable asynchronous I/O.
Other software	Korn Shell (ksh) is required. Configure the I/O completion ports (IOCP) on the operating system. To authenticate IBM Spectrum Protect users with a Lightweight Directory Access Protocol (LDAP) server, you must use one of the following directory servers: • Microsoft Active Directory (Windows Server 2012, Windows Server 2012 R2) • IBM Security Directory Server V6.3 • IBM Security Directory Server V6.4

Restriction: Raw logical volumes are not supported.

Compatibility of the IBM Spectrum Protect server with other DB2 products on the system

You can install other products that deploy and use DB2 products on the same system as the IBM Spectrum Protect Version 8.1.2 server, with some limitations.

To install and use other products that use a DB2 product on the same system as the IBM Spectrum Protect server, ensure that the following criteria are met:

Table 9. Compatibility of the IBM Spectrum Protect server with other DB2 products on the system

Criterion	Instructions
Version level	The other products that use a DB2 product must use DB2 version 9 or later. DB2 products include product encapsulation and segregation support beginning with Version 9. Starting with this version, you can run multiple copies of DB2 products, at different code levels, on the same system. For details, see the information about multiple DB2 copies in the DB2 product information.

Table 9. Compatibility of the IBM Spectrum Protect server with other DB2 products on the system (continued)

Criterion	Instructions
User IDs and directories	Ensure that the user IDs, fence user IDs, installation location, other directories, and related information are not shared across DB2 installations. Your specifications must be different from the IDs and locations that you used for the IBM Spectrum Protect server installation and configuration. If you used the **dsmicfgx** wizard to configure the server, these are values that you entered when running the wizard. If you used the manual configuration method, review the procedures that you used if necessary to recall the values that were used for the server.
Resource allocation	Consider the resources and capability of the system compared to the requirements for both the IBM Spectrum Protect server and the other applications that use the DB2 product. To provide sufficient resources for the other DB2 applications, you might have to change the IBM Spectrum Protect server settings so that the server uses less system memory and resources. Similarly, if the workloads for the other DB2 applications compete with the IBM Spectrum Protect server for processor or memory resources, the performance of the server in handling the expected client workload or other server operations might be adversely affected. To segregate resources and provide more capability for the tuning and allocation of processor, memory, and other system resources for multiple applications, consider using logical partition (LPAR), workload partition (WPAR), or other virtual workstation support. For example, run a DB2 application on its own virtualized system.

IBM Installation Manager

IBM Spectrum Protect uses IBM Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install IBM Spectrum Protect. It must remain installed on the system so that IBM Spectrum Protect can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering

An installable unit of a software product.

The IBM Spectrum Protect offering contains all of the media that IBM Installation Manager requires to install IBM Spectrum Protect.

Package

The group of software components that are required to install an offering.

The IBM Spectrum Protect package contains the following components:

- IBM Installation Manager installation program
- IBM Spectrum Protect offering

Package group

A set of packages that share a common parent directory.

The default package group for the IBM Spectrum Protect package is `IBM Installation Manager`.

Repository

A remote or local storage area for data and other application resources.

The IBM Spectrum Protect package is stored in a repository on IBM Fix Central.

Shared resources directory

A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of IBM Spectrum Protect.

Worksheets for planning details for the server

You can use the worksheets to help you plan the amount and location of storage needed for the IBM Spectrum Protect server. You can also use them to keep track of names and user IDs.

Item	Space required	Number of directories	Location of directories
The database			
Active log			
Archive log			
Optional: Log mirror for the active log			
Optional: Secondary archive log (failover location for archive log)			

Item	Names and user IDs	Location
The *instance user ID* for the server, which is the ID you use to start and run the IBM Spectrum Protect server		
The *home directory* for the server, which is the directory that contains the instance user ID		

Item	Names and user IDs	Location
The database instance name		
The *instance directory* for the server, which is a directory that contains files specifically for this server instance (the server options file and other server-specific files)		
The server name, use a unique name for each server		

Capacity planning

Capacity planning for IBM Spectrum Protect includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.

Estimating space requirements for the database

To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.

About this task

Consider using at least 25 GB for the initial database space. Provision file system space appropriately. A database size of 25 GB is adequate for a test environment or a library-manager-only environment. For a production server supporting client workloads, the database size is expected to be larger. If you use random-access disk (DISK) storage pools, more database and log storage space is needed than for sequential-access storage pools.

The maximum size of the IBM Spectrum Protect database is 6 TB.

For information about sizing the database in a production environment that is based on the number of files and on storage pool size, see the following topics.

Estimating database space requirements based on the number of files

If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.

About this task

To estimate space requirements that is based on the maximum number of files in server storage, use the following guidelines:

- 600 - 1000 bytes for each stored version of a file, including image backups.

  Restriction: The guideline does not include space that is used during data deduplication.
- 100 - 200 bytes for each cached file, copy storage pool file, active-data pool file, and deduplicated file.

- Additional space is required for database optimization to support varying data-access patterns and to support server back-end processing of the data. The amount of extra space is equal to 50% of the estimate for the total number of bytes for file objects.

In the following example for a single client, the calculations are based on the maximum values in the preceding guidelines. The examples do not take into account that you might use file aggregation. In general, when you aggregate small files, it reduces the amount of required database space. File aggregation does not affect space-managed files.

Procedure

1. Calculate the number of file versions. Add each of the following values to obtain the number of file versions:
    a. Calculate the number of backed-up files. For example, as many as 500,000 client files might be backed up at a time. In this example, storage policies are set to keep up to three copies of backed up files:

       ```
       500,000 files * 3 copies = 1,500,000 files
       ```

    b. Calculate the number of archive files. For example, as many as 100,000 client files might be archived copies.
    c. Calculate the number of space-managed files. For example, as many as 200,000 client files might be migrated from client workstations.

    Using 1000 bytes per file, the total amount of database space that is required for the files that belong to the client is 1.8 GB:

    ```
    (1,500,000 + 100,000 + 200,000) * 1000 = 1.8 GB
    ```

2. Calculate the number of cached files, copy storage-pool files, active-data pool files, and deduplicated files:
    a. Calculate the number of cached copies. For example, caching is enabled in a 5 GB disk storage pool. The high migration threshold of the pool is 90% and the low migration threshold of the pool is 70%. Thus, 20% of the disk pool, or 1 GB, is occupied by cached files.

       If the average file size is about 10 KB, approximately 100,000 files are in cache at any one time:

       ```
       100,000 files * 200 bytes = 19 MB
       ```

    b. Calculate the number of copy storage-pool files. All primary storage pools are backed up to the copy storage pool:

       ```
       (1,500,000 + 100,000 + 200,000) * 200 bytes = 343 MB
       ```

    c. Calculate the number of active storage-pool files. All the active client-backup data in primary storage pools is copied to the active-data storage pool. Assume that 500,000 versions of the 1,500,000 backup files in the primary storage pool are active:

       ```
       500,000 * 200 bytes = 95 MB
       ```

    d. Calculate the number of deduplicated files. Assume that a deduplicated storage pool contains 50,000 files:

       ```
       50,000 * 200 bytes = 10 MB
       ```

    Based on the preceding calculations, about 0.5 GB of extra database space is required for the client's cached files, copy storage-pool files, active-data pool files, and deduplicated files.

3. Calculate the amount of extra space that is required for database optimization. To provide optimal data access and management by the server, extra database space is required. The amount of extra database space is equal to 50% of the total space requirements for file objects.

   ```
   (1.8 + 0.5) * 50% = 1.2 GB
   ```
4. Calculate the total amount of database space that is required for the client. The total is approximately 3.5 GB:

   ```
   1.8 + 0.5 + 1.2 = 3.5 GB
   ```
5. Calculate the total amount of database space that is required for all clients. If the client that was used in the preceding calculations is typical and you have 500 clients, for example, you can use the following calculation to estimate the total amount of database space that is required for all clients:

   ```
   500 * 3.5 = 1.7 TB
   ```

Results

Tip: In the preceding examples, the results are estimates. The actual size of the database might differ from the estimate because of factors such as the number of directories and the length of the path and file names. Periodically monitor your database and adjust its size as necessary.

What to do next

During normal operations, the IBM Spectrum Protect server might require temporary database space. This space is needed for the following reasons:

- To hold the results of sorting or ordering that are not already being kept and optimized in the database directly. The results are temporarily held in the database for processing.
- To give administrative access to the database through one of the following methods:
  - A DB2 open database connectivity (ODBC) client
  - An Oracle Java database connectivity (JDBC) client
  - Structured Query Language (SQL) to the server from an administrative-client command line

Consider using an extra 50 GB of temporary space for every 500 GB of space for file objects and optimization. See the guidelines in the following table. In the example that is used in the preceding step, a total of 1.7 TB of database space is required for file objects and optimization for 500 clients. Based on that calculation, 200 GB is required for temporary space. The total amount of required database space is 1.9 TB.

Database size	Minimum temporary-space requirement
< 500 GB	50 GB
≥ 500 GB and < 1 TB	100 GB
≥ 1 TB and < 1.5 TB	150 GB
≥ 1.5 and < 2 TB	200 GB
≥ 2 and < 3 TB	250 - 300 GB
≥ 3 and < 4 TB	350 - 400 GB

Estimating database space requirements based on storage pool capacity

To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.

The database manager and temporary space

The IBM Spectrum Protect server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

The database manager sorts data in a specific sequence, according to the SQL statement that you issue to request the data. Depending on the workload on the server, and if there is more data than the database manager can manage, the data (that is ordered in sequence) is allocated to temporary disk space. Data is allocated to temporary disk space when there is a large result set. The database manager dynamically manages the memory that is used when data is allocated to temporary disk space.

For example, expiration processing can produce a large result set. If there is not enough system memory on the database to store the result set, some of the data is allocated to temporary disk space. During expiration processing, if a node or file space are selected that are too large to process, the database manager cannot sort the data in memory. The database manager must use temporary space to sort data.

To run database operations, consider adding more database space for the following scenarios:

- The database has a small amount of space and the server operation that requires temporary space uses the remaining free space.
- The file spaces are large, or the file spaces have an assigned policy that creates many file versions.
- The IBM Spectrum Protect server must run with limited memory. The database uses the IBM Spectrum Protect server main memory to run database operations. However, if there is insufficient memory available, the IBM Spectrum Protect server allocates temporary space on disk to the database. For example, if 10G of memory is available and database operations require 12G of memory, the database uses temporary space.
- An `out of database space` error is displayed when you deploy an IBM Spectrum Protect server. Monitor the server activity log for messages that are related to database space.

Important: Do not change the DB2 software that is installed with the IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack, of DB2 software to avoid damage to the database.

Recovery log space requirements

In IBM Spectrum Protect, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.

Active and archive log space

When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.

In IBM Spectrum Protect servers V7.1 and later, the active log can be a maximum size of 512 GB. The archive log size is limited to the size of the file system that it is installed on.

Use the following general guidelines when you estimate the size of the active log:

- The suggested starting size for the active log is 16 GB.
- Ensure that the active log is at least large enough for the amount of concurrent activity that the server typically handles. As a precaution, try to anticipate the largest amount of work that the server manages at one time. Provision the active log with extra space that can be used if needed. Consider using 20% of extra space.
- Monitor used and available active log space. Adjust the size of the active log as needed, depending upon factors such as client activity and the level of server operations.
- Ensure that the directory that holds the active log is as large as, or larger than, the size of the active log. A directory that is larger than the active log can accommodate failovers, if they occur.
- Ensure that the file system that contains the active log directory has at least 8 GB of free space for temporary log movement requirements.

The suggested starting size for the archive log is 48 GB.

The archive log directory must be large enough to contain the log files that are generated since the previous full backup. For example, if you perform a full backup of the database every day, the archive log directory must be large enough to hold the log files for all the client activity that occurs during 24 hours. To recover space, the server deletes obsolete archive log files after a full backup of the database. If the archive log directory becomes full and a directory for archive failover logs does not exist, log files remain in the active log directory. This condition can cause the active log directory to fill up and stop the server. When the server restarts, some of the existing active-log space is released.

After the server is installed, you can monitor archive log utilization and the space in the archive log directory. If the space in the archive log directory fills up, it can cause the following problems:

- The server is unable to perform full database backups. Investigate and resolve this problem.
- Other applications write to the archive log directory, exhausting the space that is required by the archive log. Do not share archive log space with other applications including other IBM Spectrum Protect servers. Ensure that each server has a separate storage location that is owned and managed by that specific server.

Example: Estimating active and archive log sizes for basic client-store operations:

Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.

To determine the sizes of the active and archive logs for basic client-store operations, use the following calculation:

```
number of clients  x  files stored during each transaction
  x  log space needed for each file
```

This calculation is used in the example in the following table.

Table 10. Basic client-store operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option `TXNGROUPMAX` is 4096.
Log space that is required for each file	3053 bytes	The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.

Active log: Suggested size	19.5 GB [1]	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.
		`(300 clients x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 bytes = 3.5 GB`
		Increase that amount by the suggested starting size of 16 GB:
		`3.5 + 16 = 19.5 GB`

Archive log: Suggested size	58.5 GB [1]	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement.
		`3.5 x 3 = 10.5 GB`
		Increase that amount by the suggested starting size of 48 GB:
		`10.5 + 48 = 58.5 GB`

Table 10. Basic client-store operations (continued)

Item	Example values	Description
[1] The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log. Monitor your logs and adjust their size if necessary.		

Example: Estimating active and archive log sizes for clients that use multiple sessions:

If the client option `RESOURCEUTILIZATION` is set to a value that is greater than the default, the concurrent workload for the server increases.

To determine the sizes of the active and archive logs when clients use multiple sessions, use the following calculation:

```
number of clients  x  sessions for each client  x  files stored
  during each transaction  x  log space needed for each file
```

This calculation is used in the example in the following table.

Table 11. Multiple client sessions

Item	Example values		Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	1000	The number of client nodes that back up, archive, or migrate files every night.
Possible sessions for each client	3	3	The setting of the client option `RESOURCEUTILIZATION` is larger than the default. Each client session runs a maximum of three sessions in parallel.
Files stored during each transaction	4096	4096	The default value of the server option `TXNGROUPMAX` is 4096.
Log space that is required for each file	3053	3053	The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.

Table 11. Multiple client sessions (continued)

Item	Example values		Description

Active log: Suggested size	26.5 GB [1]	51 GB [1]	The following calculation was used for 300 clients. One GB equals 1,073,741,824 bytes.
			`(300 clients x 3 sessions for each client x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 10.5 GB`
			Increase that amount by the suggested starting size of 16 GB:
			`10.5 + 16 = 26.5 GB`
			The following calculation was used for 1000 clients. One GB equals 1,073,741,824 bytes.
			`(1000 clients x  3 sessions for each client x 4096 files store during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 = 35 GB`
			Increase that amount by the suggested starting size of 16 GB:
			`35 + 16 = 51 GB`

Archive log: Suggested size	79.5 GB [1]	153 GB [1]	Because of the requirement to be able to store archive logs across three server-database backup cycles, the estimate for the active log is multiplied by 3:
			`10.5 x 3 = 31.5 GB`
			`35 x 3 = 105 GB`
			Increase those amounts by the suggested starting size of 48 GB:
			`31.5 + 48 = 79.5 GB`
			`105 + 48 = 153 GB`

[1] The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log. Monitor your active log and adjust its size if necessary.			

Example: Estimating active and archive log sizes for simultaneous write operations:

If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.

The log space that is required for each file increases by about 200 bytes for each copy storage pool that is used for a simultaneous write operation. In the example in the following table, data is stored to two copy storage pools in addition to a primary storage pool. The estimated log size increases by 400 bytes for each file. If you use the suggested value of 3053 bytes of log space for each file, the total number of required bytes is 3453.

This calculation is used in the example in the following table.

Table 12. Simultaneous write operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option `TXNGROUPMAX` is 4096.
Log space that is required for each file	3453 bytes	3053 bytes plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.

Active log: Suggested size	20 GB [1]	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.
		`(300 clients x 4096 files stored during each transaction x 3453 bytes for each file) ÷ 1,073,741,824 bytes = 4.0 GB`
		Increase that amount by the suggested starting size of 16 GB:
		`4 + 16 = 20 GB`

Archive log: Suggested size	60 GB [1]	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the archive log requirement:
		`4 GB  x  3  = 12 GB`
		Increase that amount by the suggested starting size of 48 GB:
		`12 + 48 = 60 GB`

[1] The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.

Monitor your logs and adjust their size if necessary.

Example: Estimating active and archive log sizes for basic client store operations and server operations:

Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.

For example, migration of files from the random-access (DISK) storage pool to a sequential-access disk (FILE) storage pool uses approximately 110 bytes of log space for each file that is migrated. For example, suppose that you have 300 backup-archive clients and each one of them backs up 100,000 files every night. The files are initially stored on DISK and then migrated to a FILE storage pool. To estimate the amount of active log space that is required for the data migration, use the following calculation. The number of clients in the calculation represents the maximum number of client nodes that back up, archive, or migrate files concurrently at any time.

```
300 clients  x 100,000 files for each client   x  110 bytes =  3.1 GB
```

Add this value to the estimate for the size of the active log that calculated for basic client store operations.

Example: Estimating active and archive log sizes under conditions of extreme variation:

Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.

Example: Estimating archive log sizes with full database backups:

The IBM Spectrum Protect server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.

For example, if a full database backup occurs once a week, the archive log space must be able to contain the information in the archive log for a full week.

The difference in archive log size for daily and full database backups is shown in the example in the following table.

Table 13. Full database backups

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option `TXNGROUPMAX` is 4096.

Table 13. Full database backups (continued)

Item	Example values	Description
Log space that is required for each file	3453 bytes	3053 bytes for each file plus 200 bytes for each copy storage pool. The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes. This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.

Active log: Suggested size	20 GB [1]	Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.
		`(300 clients x 4096 files per transaction x 3453 bytes per file) ÷ 1,073,741,824 bytes = 4.0 GB`
		Increase that amount by the suggested starting size of 16 GB:
		`4 + 16 = 20 GB`

Archive log: Suggested size with a full database backup every day	60 GB [1]	Because of the requirement to be able to store archive logs across three backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement:
		`4 GB x 3 = 12 GB`
		Increase that amount by the suggested starting size of 48 GB:
		`12 + 48 = 60 GB`

Archive log: Suggested size with a full database every week	132 GB [1]	Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. Multiply the result by the number of days between full database backups:
		`(4 GB x 3 ) x 7 = 84 GB`
		Increase that amount by the suggested starting size of 48 GB:
		`84 + 48 = 132 GB`

[1] The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested starting size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.

Monitor your logs and adjust their size if necessary.

Example: Estimating active and archive log sizes for data deduplication operations:

If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

The following factors affect requirements for active and archive log space:

The amount of deduplicated data
The effect of data deduplication on the active log and archive log space depends on the percentage of data that is eligible for deduplication. If the percentage of data that can be deduplicated is relatively high, more log space is required.

The size and number of extents
Approximately 1,500 bytes of active log space are required for each extent that is identified by a duplicate-identification process. For example, if 250,000 extents are identified by a duplicate-identification process, the estimated size of the active log is 358 MB:

```
250,000 extents identified during each process x 1,500 bytes
 for each extent = 358 MB
```

Consider the following scenario. Three hundred backup-archive clients back up 100,000 files each night. This activity creates a workload of 30,000,000 files. The average number of extents for each file is two. Therefore, the total number of extents is 60,000,000, and the space requirement for the archive log is 84 GB:

```
60,000,000 extents x 1,500 bytes for each extent = 84 GB
```

A duplicate-identification process operates on aggregates of files. An aggregate consists of files that are stored in a given transaction, as specified by the `TXNGROUPMAX` server option. Suppose that the `TXNGROUPMAX` server option is set to the default of 4096. If the average number of extents for each file is two, the total number of extents in each aggregate is 8192, and the space required for the active log is 12 MB:

```
8192 extents in each aggregate x 1500 bytes for each extent =
   12 MB
```

The timing and number of the duplicate-identification processes
The timing and number of duplicate-identification processes also affects the size of the active log. Using the 12 MB active-log size that was calculated in the preceding example, the concurrent load on the active log is 120 MB if 10 duplicate-identification processes are running in parallel:

```
12 MB for each process x 10 processes = 120 MB
```

File size
Large files that are processed for duplicate identification can also affect the size of the active log. For example, suppose that a backup-archive client backs up an 80 GB, file-system image. This object can have a high number of duplicate extents if, for example, the files included in the file system image were backed up incrementally. For example, assume that a file system image has 1.2 million duplicate extents. The 1.2 million extents in this large file represent a single transaction for a duplicate-identification process. The total space in the active log that is required for this single object is 1.7 GB:

```
1,200,000 extents x 1,500 bytes for each extent = 1.7 GB
```

If other, smaller duplicate-identification processes occur at the same time as the duplicate-identification process for a single large object, the active log might not have enough space. For example, suppose that a storage pool is enabled for deduplication. The storage pool has a mixture of data, including many relatively small files that range from 10 KB to several hundred KB. The storage pool also has few large objects that have a high percentage of duplicate extents.

To take into account not only space requirements but also the timing and duration of concurrent transactions, increase the estimated size of the active log by a factor of two. For example, suppose that your calculations for space requirements are 25 GB (23.3 GB + 1.7 GB for deduplication of a large object). If deduplication processes are running concurrently, the suggested size of the active log is 50 GB. The suggested size of the archive log is 150 GB.

The examples in the following tables show calculations for active and archive logs. The example in the first table uses an average size of 700 KB for extents. The example in the second table uses an average size of 256 KB. As the examples show, the average deduplicate-extent size of 256 KB indicates a larger estimated size for the active log. To minimize or prevent operational problems for the server, use 256 KB to estimate the size of the active log in your production environment.

Table 14. Average duplicate-extent size of 700 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.
Average size of extents	700 KB	700 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average size for extents.
Extents for a given file	1,198,372 bits	6,135,667 bits	Using the average extent size (700 KB), these calculations represent the total number of extents for a given object. The following calculation was used for an 800 GB object: `(800 GB ÷ 700 KB) = 1,198,372 bits` The following calculation was used for a 4 TB object: `(4 TB ÷ 700 KB) = 6,135,667 bits`
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	1.7 GB	8.6 GB	The estimated active log space that are needed for this transaction.

Table 14. Average duplicate-extent size of 700 KB (continued)

Item	Example values		Description

Active log: Suggested total size	66 GB [1]	79.8 GB [1]	After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of two. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.
			The following calculation was used for multiple transactions and an 800 GB object:
			`(23.3 GB + 1.7 GB) x 2 = 50 GB`
			Increase that amount by the suggested starting size of 16 GB:
			`50 + 16 = 66 GB`
			The following calculation was used for multiple transactions and a 4 TB object:
			`(23.3 GB + 8.6 GB) x 2 = 63.8 GB`
			Increase that amount by the suggested starting size of 16 GB:
			`63.8 + 16 = 79.8 GB`

Archive log: Suggested size	198 GB [1]	239.4 GB [1]	Multiply the estimated size of the active log by a factor of 3.
			The following calculation was used for multiple transactions and an 800 GB object:
			`50 GB x 3 = 150 GB`
			Increase that amount by the suggested starting size of 48 GB:
			`150 + 48 = 198 GB`
			The following calculation was used for multiple transactions and a 4 TB object:
			`63.8 GB x 3 = 191.4 GB`
			Increase that amount by the suggested starting size of 48 GB:
			`191.4 + 48 = 239.4 GB`

[1] The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log. Monitor your logs and adjust their size if necessary.			

Table 15. Average duplicate-extent size of 256 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.

Table 15. Average duplicate-extent size of 256 KB (continued)

Item	Example values		Description
Average size of extents	256 KB	256 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average extent size.

Extents for a given file	3,276,800 bits	16,777,216 bits	Using the average extent size, these calculations represent the total number of extents for a given object.
			The following calculation was used for multiple transactions and an 800 GB object:
			`(800 GB ÷ 256 KB) = 3,276,800 bits`
			The following calculation was used for multiple transactions and a 4 TB object:
			`(4 TB ÷ 256 KB) = 16,777,216 bits`

Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	4.5 GB	23.4 GB	The estimated size of the active log space that is required for this transaction.

Active log: Suggested total size	71.6 GB [1]	109.4 GB [1]	After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of 2. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.
			The following calculation was used for multiple transactions and an 800 GB object:
			`(23.3 GB + 4.5 GB) x 2 = 55.6 GB`
			Increase that amount by the suggested starting size of 16 GB:
			`55.6 + 16 = 71.6 GB`
			The following calculation was used for multiple transactions and a 4 TB object:
			`(23.3 GB + 23.4 GB) x 2 = 93.4 GB`
			Increase that amount by the suggested starting size of 16 GB:
			`93.4 + 16 = 109.4 GB`

Table 15. Average duplicate-extent size of 256 KB (continued)

Item	Example values		Description

Archive log: Suggested size	214.8 GB [1]	328.2 GB [1]	The estimated size of the active log multiplied by a factor of 3.
			The following calculation was used for an 800 GB object:
			`55.6 GB x 3 = 166.8 GB`
			Increase that amount by the suggested starting size of 48 GB:
			`166.8 + 48 = 214.8 GB`
			The following calculation was used for a 4 TB object:
			`93.4 GB x 3 = 280.2 GB`
			Increase that amount by the suggested starting size of 48 GB:
			`280.2 + 48 = 328.2 GB`

[1] The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log. Monitor your logs and adjust their size if necessary.			

Active-log mirror space

The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.

Creating a log mirror is a suggested option. If you increase the size of the active log, the log mirror size is increased automatically. Mirroring the log can affect performance because of the doubled I/O activity that is required to maintain the mirror. The additional space that the log mirror requires is another factor to consider when deciding whether to create a log mirror.

If the mirror log directory becomes full, the server issues error messages to the activity log and to the `db2diag.log`. Server activity continues.

Archive-failover log space

The archive failover log is used by the server if the archive log directory runs out of space.

Specifying an archive failover log directory can prevent problems that occur if the archive log runs out of space. If both the archive log directory and the drive or file system where the archive failover log directory is located become full, the data remains in the active log directory. This condition can cause the active log to fill up, which causes the server to halt.

Monitoring space utilization for the database and recovery logs

To determine the amount of used and available active log space, you issue the **`QUERY LOG`** command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.

Active log

If the amount of available active log space is too low, the following messages are displayed in the activity log:

ANR4531I: IC_AUTOBACKUP_LOG_USED_SINCE_LAST_BACKUP_TRIGGER

This message is displayed when the active log space exceeds the maximum specified size. The IBM Spectrum Protect server starts a full database backup.

To change the maximum log size, halt the server. Open the `dsmserv.opt` file, and specify a new value for the `ACTIVELOGSIZE` option. When you are finished, restart the server.

ANR0297I: IC_BACKUP_NEEDED_LOG_USED_SINCE_LAST_BACKUP

This message is displayed when the active log space exceeds the maximum specified size. You must back up the database manually.

To change the maximum log size, halt the server. Open the `dsmserv.opt` file, and specify a new value for the `ACTIVELOGSIZE` option. When you are finished, restart the server.

ANR4529I: IC_AUTOBACKUP_LOG_UTILIZATION_TRIGGER

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. If at least one full database backup has occurred, the IBM Spectrum Protect server starts an incremental database backup. Otherwise, the server starts a full database backup.

ANR0295I: IC_BACKUP_NEEDED_LOG_UTILIZATION

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. You must back up the database manually.

Archive log

If the amount of available archive log space is too low, the following message is displayed in the activity log:

ANR0299I: IC_BACKUP_NEEDED_ARCHLOG_USED

The ratio of used archive-log space to available archive-log space exceeds the log utilization threshold. The IBM Spectrum Protect server starts a full automatic database backup.

Database

If the amount of space available for database activities is too low, the following messages are displayed in the activity log:

ANR2992W: IC_LOG_FILE_SYSTEM_UTILIZATION_WARNING_2

The used database space exceeds the threshold for database space utilization. To increase the space for the database, use the **`EXTEND DBSPACE`** command, the **`EXTEND DBSPACE`** command, or the DSMSERV FORMAT utility with the **`DBDIR`** parameter.

ANR1546W: FILESYSTEM_DBPATH_LESS_1GB

The available space in the directory where the server database files are located is less than 1 GB.

When an IBM Spectrum Protect server is created with the DSMSERV FORMAT utility or with the configuration wizard, a server database and recovery log are also created. In addition, files are created to hold database information used by the database manager. The path specified in this message indicates the location of the database information used by the database manager. If space is unavailable in the path, the server can no longer function.

You must add space to the file system or make space available on the file system or disk.

Deleting installation rollback files

You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

About this task

To delete the files that are no longer needed, use either the installation graphical wizard or the command line in console mode.

Deleting installation rollback files by using a graphical wizard

You can delete certain installation files that were saved during installation process by using the IBM Installation Manager user interface.

Procedure

1. Open IBM Installation Manager.

   In the directory where IBM Installation Manager is installed, go to the `eclipse` subdirectory (for example, `/opt/IBM/InstallationManager/eclipse`), and issue the following command to start IBM Installation Manager:

   ```
   ./IBMIM
   ```
2. Click **File** > **Preferences**.
3. Select **Files for Rollback**.
4. Click **Delete Saved Files** and click **OK**.

Deleting installation rollback files by using the command line

You can delete certain installation files that were saved during the installation process by using the command line.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:

   ```
   eclipse/tools
   ```

   For example:

   ```
   /opt/IBM/InstallationManager/eclipse/tools
   ```
2. From the `tools` directory, issue the following command to start an IBM Installation Manager command line:

   ```
   ./imcl -c
   ```
3. Enter `P` to select Preferences.

4. Enter `3` to select Files for Rollback.
5. Enter `D` to Delete the Files for Rollback.
6. Enter `A` to Apply Changes and Return to Preferences Menu.
7. Enter `C` to leave the Preference Menu.
8. Enter `X` to Exit Installation Manager.

Server naming best practices

Use these descriptions as a reference when you install or upgrade an IBM Spectrum Protect server.

Instance user ID

The instance user ID is used as the basis for other names related to the server instance. The instance user ID is also called the instance owner.

For example: tsminst1

The instance user ID is the user ID that must have ownership or read/write access authority to all directories that you create for the database and the recovery log. The standard way to run the server is under the instance user ID. That user ID must also have read/write access to the directories that are used for any **`FILE`** device classes.

Home directory for the instance user ID

The home directory can be created when creating the instance user ID, by using the option (`-m`) to create a home directory if it does not exist already. Depending on local settings, the home directory might have the form: `/home/`*`instance_user_ID`*

For example: `/home/tsminst1`

The home directory is primarily used to contain the profile for the user ID and for security settings.

Database instance name

The database instance name must be the same as the instance user ID under which you run the server instance.

For example: `tsminst1`

Instance directory

The instance directory is a directory that contains files specifically for a server instance (the server options file and other server-specific files). It can have any name that you want. For easier identification, use a name that ties the directory to the instance name.

You can create the instance directory as a subdirectory of the home directory for the instance user ID. For example: `/home/`*`instance_user_ID/instance_user_ID`*

The following example places the instance directory in the home directory for user ID tsminst1: `/home/tsminst1/tsminst1`

You can also create the directory in another location, for example: `/tsmserver/tsminst1`

The instance directory stores the following files for the server instance:

- The server options file, `dsmserv.opt`
- The server key database file, `cert.kdb`, and the `.arm` files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
- Device configuration file, if the `DEVCONFIG` server option does not specify a fully qualified name
- Volume history file, if the `VOLUMEHISTORY` server option does not specify a fully qualified name
- Volumes for **DEVTYPE=FILE** storage pools, if the directory for the device class is not fully specified, or not fully qualified
- User exits
- Trace output (if not fully qualified)

Database name

The database name is always TSMDB1, for every server instance. This name cannot be changed.

Server name

The server name is an internal name for IBM Spectrum Protect, and is used for operations that involve communication among multiple IBM Spectrum Protect servers. Examples include server-to-server communication and library sharing.

The server name is also used when you add the server to the Operations Center so that it can be managed using that interface. Use a unique name for each server. For easy identification in the Operations Center (or from a **`QUERY SERVER`** command), use a name that reflects the location or purpose of the server. Do not change the name of an IBM Spectrum Protect server after it is configured as a hub or spoke server.

If you use the wizard, the default name that is suggested is the host name of the system that you are using. You can use a different name that is meaningful in your environment. If you have more than one server on the system and you use the wizard, you can use the default name for only one of the servers. You must enter a unique name for each server.

For example:

```
PAYROLL
SALES
```

Directories for database space and recovery log

The directories can be named according to local practices. For easier identification, consider using names that tie the directories to the server instance.

For example, for the archive log:

```
/tsminst1_archlog
```

Installation directories

Installation directories for the IBM Spectrum Protect server include the server, DB2, device, language, and other directories. Each one contains several additional directories.

The (`/opt/tivoli/tsm/server/bin`) is the default directory that contains server code and licensing.

The DB2 product that is installed as part of the installation of the IBM Spectrum Protect server has the directory structure as documented in DB2 information sources. Protect these directories and files as you do the server directories. The default directory is `/opt/tivoli/tsm/db2`.

You can use US English, German, French, Italian, Spanish, Brazilian Portuguese, Korean, Japanese, traditional Chinese, simplified Chinese, Chinese GBK, Chinese Big5, and Russian.

Chapter 2. Installing the server components

To install the Version 8.1.2 server components, you can use the installation wizard, the command line in console mode, or silent mode.

About this task

Using the IBM Spectrum Protect installation software, you can install the following components:

- server

  Tip: The database (DB2), the Global Security Kit (GSKit) and IBM Java Runtime Environment (JRE) are automatically installed when you select the server component.
- server languages
- license
- devices
- IBM Spectrum Protect for SAN
- Operations Center

Allow approximately 30 - 45 minutes to install a V8.1.2 server, using this guide.

Obtaining 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

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

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 appropriate package file from one of the following websites.
   - Download the server package from Passport Advantage or Fix Central.
   - For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.
2. If you downloaded the package from an IBM download site, 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 4042944
      - IBM Spectrum Protect Extended Edition technote 4042945

  - IBM Spectrum Protect for Data Retention technote 4042946

b. Download the package file 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 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.bin
```

d. Extract the package by issuing the following command:

```
./package_name.bin
```

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

```
8.1.x.000-IBM-SPSRV-Linuxx86_64.bin
8.1.x.000-IBM-SPSRV-Linuxs390x.bin
8.1.x.000-IBM-SPSRV-Linuxppc64le.bin
```

3. Select one of the following methods of installing IBM Spectrum Protect:
   - “Installing IBM Spectrum Protect by using the installation wizard”
   - “Installing IBM Spectrum Protect by using console mode” on page 55
   - “Installing IBM Spectrum Protect in silent mode” on page 56
4. After you install IBM Spectrum Protect, and before you customize it for your use, go to the IBM Support Portal. Click **Support and downloads** and apply any applicable fixes.

Installing IBM Spectrum Protect by using the installation wizard

You can install the server by using the IBM Installation Manager graphical wizard.

Before you begin

Take the following actions before you start the installation:

- 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.

Procedure

Install IBM Spectrum Protect by using this method:

Option	Description
Installing the software from a downloaded package:	1. Change to the directory where you downloaded the package. 2. Start the installation wizard by issuing the following command: `./install.sh`

What to do next

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

You can view installation log files by clicking **File** > **View Log** from the Installation Manager tool. To collect these log files, click **Help** > **Export Data for Problem Analysis** from the Installation Manager tool.

- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
- After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

Installing IBM Spectrum Protect by using console mode

You can install IBM Spectrum Protect by using the command line in console mode.

Before you begin

Take the following actions before you start the installation:

- 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.

Procedure

Install IBM Spectrum Protect by using this method:

Option	Description
Installing the software from a downloaded package:	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 panel, specify `G` to generate the responses.

What to do next

- 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
  ```
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
- After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

Installing IBM Spectrum Protect in silent mode

You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the `input` directory where the installation package is extracted:

`install_response_sample.xml`
Use this file to install the IBM Spectrum Protect components.

`update_response_sample.xml`
Use this file to upgrade the IBM Spectrum Protect components.

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.
2. If you install the server and Operations Center in silent mode, create a password for the Operations Center truststore in the response file.

   If you are using the `install_response_sample.xml` file, add the password in the following line of the file, where *mypassword* represents the password:

   ```
   <variable name='ssl.password' value='mypassword' />
   ```

   For more information about this password, see Installation checklist

   Tip: To upgrade the Operations Center, the truststore password is not required if you are using the `update_response_sample.xml` file.
3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value *response_file* represents the response file path and file name:
   - `./install.sh -s -input response_file -acceptLicense`

What to do next

- 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
  ```
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
- After you install a new server, review Taking the first steps after you install IBM Spectrum Protect to learn about configuring your server.

Installing server language packages

Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Before you begin

For instructions on installing storage agent language packages, see Language pack configuration for storage agents.

Server language locales

Use either the default language package option or select another language package to display server messages and help.

This language package is automatically installed for the following default language option for IBM Spectrum Protect server messages and help:

- LANGUAGE en_US

For languages or locales other than the default, install the language package that your installation requires.

You can use the languages that are shown:

Table 16. Server languages for Linux

LANGUAGE	LANGUAGE option value
Chinese, Simplified	zh_CN
	zh_CN.gb18030
	zh_CN.utf8
Chinese, Traditional	Big5 / Zh_TW
	zh_TW
	zh_TW.utf8
English, United States	en_US
	en_US.utf8
French	fr_FR
	fr_FR.utf8
German	de_DE
	de_DE.utf8
Italian	it_IT
	it_IT.utf8
Japanese	ja_JP
	ja_JP.utf8
Korean	ko_KR
	ko_KR.utf8
Portuguese, Brazilian	pt_BR
	pt_BR.utf8

Table 16. Server languages for Linux (continued)

LANGUAGE	LANGUAGE option value
Russian	ru_RU
	ru_RU.utf8
Spanish	es_ES
	es_ES.utf8

Restriction: For Operations Center users, some characters might not be displayed properly if the web browser does not use the same language as the server. If this problem occurs, set the browser to use the same language as the server.

Configuring a language package

After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with IBM Spectrum Protect.

About this task

To set support for a certain locale, complete one of the following tasks:

- Set the `LANGUAGE` option in the server options file to the name of the locale that you want to use. For example:

  To use the `it_IT` locale, set the `LANGUAGE` option to `it_IT`. See "Server language locales" on page 57.
- If you are starting the server in the foreground, set the `LC_ALL` environment variable to match the value that is set in the server options file. For example, to set the environment variable for Italian, enter the following value:

  ```
  export LC_ALL=it_IT
  ```

If the locale is successfully initialized, it formats the date, time, and number for the server. If the locale is not successfully initialized, the server uses the US English message files and the date, time, and number format.

Updating a language package

You can modify or update a language package by using the IBM Installation Manager.

About this task

You can install another language package within the same IBM Spectrum Protect instance.

- Use the **Modify** function of IBM Installation Manager to install another language package.
- Use the **Update** function of IBM Installation Manager to update to newer versions of the language packages.

Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.

Chapter 3. Taking the first steps after you install IBM Spectrum Protect

After you install Version 8.1.2, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the IBM Spectrum Protect instance.

About this task

1. Update the kernel parameter values.

   See "Tuning kernel parameters."
2. Create the directories and user ID for the server instance. See "Creating the user ID and directories for the server instance" on page 61.
3. Configure a server instance. Select one of the following options:
   - Use the configuration wizard, the preferred method. See "Configuring IBM Spectrum Protect by using the configuration wizard" on page 63.
   - Manually configure the new instance. See "Configuring the server instance manually" on page 63. Complete the following steps during a manual configuration.
     a. Set up your directories and create the IBM Spectrum Protect instance. See "Creating the server instance" on page 64.
     b. Create a new server options file by copying the sample file to set up communications between the server and clients. See "Configuring server and client communications" on page 65.
     c. Issue the **DSMSERV FORMAT** command to format the database. See "Formatting the database and log" on page 69.
     d. Configure your system for database backup. See "Preparing the database manager for database backup" on page 70.
4. Configure options to control when database reorganization runs. See "Configuring server options for server database maintenance" on page 72.
5. Start the server instance if it is not already started.

   See "Starting the server instance" on page 73.
6. Register your license. See "Registering licenses" on page 79.
7. Prepare your system for database backups. See "Specifying a device class in preparation for database backups" on page 79.
8. Monitor the server. See "Monitoring the server" on page 80.

Tuning kernel parameters

For IBM Spectrum Protect and DB2 to install and operate correctly on Linux, you must update the kernel configuration parameters.

About this task

If you do not update these parameters, the installation of DB2 and IBM Spectrum Protect might fail. Even if installation is successful, operational problems might occur if you do not set parameter values.

Updating kernel parameters

DB2 automatically increases interprocess communication (IPC) kernel parameter values to the preferred settings.

About this task

To update the kernel parameters on Linux servers, complete the following steps:

Procedure

1. Issue the **`ipcs -l`** command to list the parameter values.
2. Analyze the results to determine whether any changes are required for your system. If changes are required, you can set the parameter in the `/etc/sysctl.conf` file. The parameter value is applied when the system starts.

What to do next

For Red Hat Enterprise Linux 6 (RHEL6), you must set the `kernel.shmmax` parameter in the `/etc/sysctl.conf` file before automatically starting the IBM Spectrum Protect server on system startup.

For details about the DB2 database for Linux, see the DB2 product information.

Suggested values

Ensure that the values for kernel parameters are sufficient to prevent operational problems from occurring when you run the IBM Spectrum Protect server.

About this task

The following table contains the suggested kernel parameter settings to run both IBM Spectrum Protect and DB2.

Parameter	Description	Preferred value
kernel.randomize_va_space	The **`kernel.randomize_va_space`** parameter configures the kernels use of memory ASLR. When you set the value to 0, `kernel.randomize_va_space=0`, it disables ASLR. DB2 data servers rely on fixed addresses for certain shared memory objects, and the ASLR can cause errors for some activities. To learn more details about the Linux ASLR and DB2, see the technote at: http://www.ibm.com/support/docview.wss?uid=swg21365583.	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.	0
vm.overcommit_memory	The **`vm.overcommit_memory`** parameter influences how much virtual memory the kernel can permit be allocated. For more information about kernel parameters, see the DB2 product information.	0

Creating the user ID and directories for the server instance

Create the user ID for the IBM Spectrum Protect server instance and create the directories that the server instance needs for database and recovery logs.

Before you begin

Review the information about planning space for the server before you complete this task. See "Worksheets for planning details for the server" on page 30.

Procedure

1. Create the user ID that will own the server instance. You use this user ID when you create the server instance in a later step.

   Create a user ID and group that will be the owner of the server instance.

   a. The following commands can be run from an administrative user ID that will set up the user and group. Create the user ID and group in the home directory of the user.

      Restriction: In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID and group name must comply with the following rules:

      - The length must be 8 characters or less.
      - The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
      - The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

      For example, create user ID `tsminst1` in group `tsmsrvrs`. The following examples show how to create this user ID and group using operating system commands.

      ```
      groupadd tsmsrvrs -g 1111
      useradd -d /home/tsminst1 -u 2222 -g 1111 -s /bin/bash tsminst1
      passwd tsminst1
      ```

      Restriction: DB2 does not support direct operating system user authentication through LDAP.

   b. Log off, then log in to your system. Change to the user account that you just created. Use an interactive login program, such as telnet, so that you are prompted for the password and can change it if necessary.

2. Create directories that the server requires.

Create empty directories for each item in the table and ensure that the directories are owned by the new user ID you just created. Mount the associated storage to each directory for the active log, archive log, and database directories.

Item	Example commands for creating the directories	Your directories
The *instance directory* for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	`mkdir /tsminst1`	

The database directories	`mkdir /tsmdb001`
	`mkdir /tsmdb002`
	`mkdir /tsmdb003`
	`mkdir /tsmdb004`	

Active log directory	`mkdir /tsmlog`	
Archive log directory	`mkdir /tsmarchlog`	
Optional: Directory for the log mirror for the active log	`mkdir /tsmlogmirror`	
Optional: Secondary archive log directory (failover location for archive log)	`mkdir /tsmarchlogfailover`	

When a server is initially created by using the **DSMSERV FORMAT** utility or the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

3. Log off the new user ID.

Configuring the IBM Spectrum Protect server

After you have installed the server and prepared for the configuration, configure the server instance.

About this task

Configure an IBM Spectrum Protect server instance by selecting one of the following options:

- Use the IBM Spectrum Protect configuration wizard on your local system. See "Configuring IBM Spectrum Protect by using the configuration wizard" on page 63.
- Manually configure the new IBM Spectrum Protect instance. See "Configuring the server instance manually" on page 63. Complete the following steps during a manual configuration.
  1. Set up the directories and create the IBM Spectrum Protect instance. See "Creating the server instance" on page 64.
  2. Create a new server options file by copying the sample file in order to set up communications between the IBM Spectrum Protect server and clients. See "Configuring server and client communications" on page 65 .

3. Issue the DSMSERV FORMAT command to format the database. See “Formatting the database and log” on page 69.
4. Configure your system for database backup. See “Preparing the database manager for database backup” on page 70.

Configuring IBM Spectrum Protect by using the configuration wizard

The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the IBM Spectrum Protect server program.

Before you begin

Before you begin to use the configuration wizard, you must complete all preceding steps to prepare for the configuration. These steps include installing IBM Spectrum Protect, creating the database and log directories, and creating the directories and user ID for the server instance.

Procedure

1. Ensure that the following requirements are met:
   - The system where you installed IBM Spectrum Protect must have the X Window System client. You must also be running an X Window System server on your desktop.
   - The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in the `/etc/ssh/` directory. Also, ensure that the SSH daemon service has access rights for connecting to the system by using the `localhost` value.
   - You must be able to log in to IBM Spectrum Protect with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
   - Restart the server before you proceed with the Configuration wizard.
2. Start the local version of the wizard:

   > Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory. This wizard can be only run as a root user.

   Follow the instructions to complete the configuration. The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Configuring the server instance manually

After installing IBM Spectrum Protect Version 8.1.2, you can configure IBM Spectrum Protect manually instead of using the configuration wizard.

Creating the server instance

Create an IBM Spectrum Protect instance by issuing the **`db2icrt`** command.

About this task

You can have one or more server instances on one workstation.

Important: Before you run the **`db2icrt`** command, verify the following items:

- The home directory for the user (`/home/tsminst1`) exists. If there is no home directory, you must create it.

  The instance directory stores the following core files that are generated by the IBM Spectrum Protect server:
  - The server options file, `dsmserv.opt`
  - The server key database file, `cert.kdb`, and the `.arm` files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
  - Device configuration file, if the `DEVCONFIG` server option does not specify a fully qualified name
  - Volume history file, if the `VOLUMEHISTORY` server option does not specify a fully qualified name
  - Volumes for **DEVTYPE=FILE** storage pools, if the directory for the device class is not fully specified, or not fully qualified
  - User exits
  - Trace output (if not fully qualified)
- A shell configuration file (for example, `.profile`) exists in the home directory. The root user and instance-user ID must have write permission to this file. For more information, see the DB2 product information. Search for Linux and UNIX environment variable settings.

1. Log in using the root user ID and create an IBM Spectrum Protect instance. The name of the instance must be the same name as the user that owns the instance. Use the **`db2icrt`** command and enter the command on one line:

   ```
   /opt/tivoli/tsm/db2/instance/db2icrt -a server -u
   instance_name instance_name
   ```

   For example, if your user ID for this instance is tsminst1, use the following command to create the instance. Enter the command on one line.

   ```
   /opt/tivoli/tsm/db2/instance/db2icrt -a server -u
   tsminst1 tsminst1
   ```

   Remember: From this point on, use this new user ID when you configure your IBM Spectrum Protect server. Log out of the root user ID and log in under the new instance-user ID.
2. Change the default directory for the database to be the same as the instance directory for the server. If you have multiple servers, log in under the instance ID for each server. Issue this command:

   ```
   db2 update dbm cfg using dftdbpath instance_directory
   ```

   For example, where instance_directory is the instance user ID:

   ```
   db2 update dbm cfg using dftdbpath /tsminst1
   ```
3. Modify the library path to use the version of the IBM Global Security Kit (GSKit) that is installed with the server. In the following examples, *server_bin_directory* is a subdirectory of the server installation directory. For example, `/opt/tivoli/tsm/server/bin`.

- You must update the following files to set the library path when DB2 or the server are started:

  Bash or Korn shell example:

  `instance_users_home_directory``/sqllib/userprofile`

  C shell example:

  `instance_users_home_directory``/sqllib/usercshrc`
- Add the following entry to the *`instance_users_home_directory`*`/sqllib/userprofile` (Bash or Korn shell) file. Each entry is on one line.

  ```
  LD_LIBRARY_PATH=server_bin_directory/dbbkapi:
  /opt/ibm/lib:/opt/ibm/lib64:/usr/lib64:$LD_LIBRARY_PATH

  export LD_LIBRARY_PATH
  ```

  Remember: The following entries must be in the library path:
  - `/usr/local/ibm/gsk8_64/lib64`
  - `/opt/ibm/lib`
  - `/opt/ibm/lib64`
  - `/usr/lib64`
- Add the following entry to the *`instance_users_home_directory`*`/sqllib/usercshrc` (C shell) file, on one line:

  ```
  setenv LD_LIBRARY_PATH server_bin_directory/dbbkapi:
  /opt/ibm/lib:/opt/ibm/lib64:/usr/lib64:$LD_LIBRARY_PATH
  ```
- Verify the library path settings and that the GSKit is version 8.0.14.43 or later. Issue the following commands:

  ```
  echo $LD_LIBRARY_PATH
  gsk8capicmd_64 -version
  gsk8ver_64
  ```

  If your GSKit version is not 8.0.14.43 or later, you must reinstall the IBM Spectrum Protect server. The reinstallation ensures that the correct GSKit version is available.

4. Create a new server options file. See “Configuring server and client communications.”

Configuring server and client communications

A default sample server options file, `dsmserv.opt.smp`, is created during IBM Spectrum Protect installation in the `/opt/tivoli/tsm/server/bin` directory. You must set up communications between the server and clients by creating a new server options file. To do so, copy the sample file to the directory for the server instance.

About this task

Ensure that you have a server instance directory, for example `/tsminst1`, and copy the sample file to this directory. Name the new file `dsmserv.opt` and edit the options. Complete this set-up before you initialize the server database. Each sample or default entry in the sample options file is a comment, a line beginning with an asterisk (*). Options are not case-sensitive and one or more blank spaces are allowed between keywords and values.

When editing the options file, follow these guidelines:

- Remove the asterisk at the beginning of the line to activate an option.
- Begin entering the options in any column.
- Enter only one option per line, and the option must be on only one line.

- If you make multiple entries for a keyword, the IBM Spectrum Protect server uses the last entry.

If you change the server options file, you must restart the server for the changes to take effect.

You can specify one or more of the following communication methods:
- TCP/IP Version 4 or Version 6
- Shared memory
- Secure Sockets Layer (SSL)

Tip: You can authenticate passwords with the LDAP directory server, or authenticate passwords with the IBM Spectrum Protect server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

Setting TCP/IP options:

Select from a range of TCP/IP options for the IBM Spectrum Protect server or retain the default.

About this task

The following is an example of a list of TCP/IP options that you can use to set up your system.

```
commmethod        tcpip
  tcpport         1500
  tcpwindowsize   0
  tcpnodelay      yes
```

Tip: You can use TCP/IP Version 4, Version 6, or both.

TCPPORT
The server port address for TCP/IP and SSL communication. The default value is 1500.

TCPWINDOWSIZE
Specifies the size of the TCP/IP buffer that is used when sending or receiving data. The window size that is used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but can improve performance.

You can specify an integer from 0 to 2048. To use the default window size for the operating system, specify 0.

TCPNODELAY
Specifies whether or not the server sends small messages or lets TCP/IP buffer the messages. Sending small messages can improve throughput but increases the number of packets sent over the network. Specify YES to send small messages or NO to let TCP/IP buffer them. The default is YES.

TCPADMINPORT
Specifies the port number on which the server TCP/IP communication driver is to wait for TCP/IP or SSL-enabled communication requests other than client sessions. The default is the value of `TCPPORT`.

SSLTCPPORT

(SSL-only) Specifies the Secure Sockets Layer (SSL) port number on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line backup-archive client and the command-line administrative client.

SSLTCPADMINPORT

(SSL-only) Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client.

Setting shared memory options:

You can use shared memory communications between clients and servers on the same system. To use shared memory, TCP/IP Version 4 must be installed on the system.

About this task

The following example shows a shared memory setting:

```
commmethod     sharedmem
   shmport     1510
```

In this example, **SHMPORT** specifies the TCP/IP port address of a server when using shared memory. Use the **SHMPORT** option to specify a different TCP/IP port. The default port address is 1510.

COMMMETHOD can be used multiple times in the IBM Spectrum Protect server options file, with a different value each time. For example, the following example is possible:

```
commmethod tcpip
commmethod sharedmem
```

You might receive the following message from the server when using shared memory:

```
ANR9999D shmcomm.c(1598): ThreadId<39>
Error from msgget (2), errno = 28
```

The message means that a message queue must be created but the system limit for the maximum number of message queues (**MSGMNI**) would be exceeded.

To find out the maximum number of message queues (**MSGMNI**) on your system, issue the following command:

```
cat /proc/sys/kernel/msgmni
```

To increase the **MSGMNI** value on your system, issue the following command:

```
sysctl -w kernel.msgmni=n
```

where **n** is the maximum number of message queues that you want the system to allow.

Setting Secure Sockets Layer options:

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Before you begin

SSL is the standard technology for creating encrypted sessions between servers and clients. SSL provides a secure channel for servers and clients to communicate over open communication paths. With SSL, the identity of the server is verified through the use of digital certificates.

To ensure better system performance, use SSL only for sessions when it is needed. Consider adding additional processor resources on the IBM Spectrum Protect server to manage the increased requirements.

Formatting the database and log

Use the **DSMSERV FORMAT** utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.

After you set up server communications, you are ready to initialize the database. Ensure that you log in by using the instance user ID. Do not place the directories on file systems that might run out of space. If certain directories (for example, the archive log) become unavailable or full, the server stops. See Capacity planning for more details.

For optimal performance and to facilitate I/O, specify at least two equally sized containers or Logical Unit Numbers (LUNs) for the database. In addition, each active log and archive log needs its own container or LUN.

Setting the exit list handler

Set the **DB2NOEXITLIST** registry variable to `ON` for each server instance. Log on to the system as the server instance owner and issue this command:

```
db2set -i server_instance_name DB2NOEXITLIST=ON
```

For example:

```
db2set -i tsminst1 DB2NOEXITLIST=ON
```

Initializing a server instance

Use the **DSMSERV FORMAT** utility to initialize a server instance. For example, if the server instance directory is */tsminst1*, issue the following commands:

```
cd /tsminst1
dsmserv format dbdir=/tsmdb001 activelogsize=32768
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

Tip: If DB2 does not start after you issue the **DSMSERV FORMAT** command, you might need to disable the file system mount option `NOSUID`. If this option is set on the file system that contains the DB2 instance owner directory, or on any file system that contains the DB2 database, active logs, archive logs, failover logs, or mirrored logs, the option must be disabled to start the system.

After you disable the `NOSUID` option, remount the file system and then start DB2 by issuing the following command:

```
db2start
```

Related information:

DSMSERV FORMAT (Format the database and log)

Preparing the database manager for database backup

To back up the data in the database to IBM Spectrum Protect, you must enable the database manager and configure the IBM Spectrum Protect application programming interface (API).

About this task

Starting with IBM Spectrum Protect V7.1, it is no longer necessary to set the API password during a manual configuration of the server. If you set the API password during the manual configuration process, attempts to back up the database might fail.

If you use the configuration wizard to create an IBM Spectrum Protect server instance, you do not have to complete these steps. If you are configuring an instance manually, complete the following steps before you issue either the **BACKUP DB** or the **RESTORE DB** commands.

Attention: If the database is unusable, the entire IBM Spectrum Protect server is unavailable. If a database is lost and cannot be recovered, it might be difficult or impossible to recover data that is managed by that server. Therefore, it is critically important to back up the database.

In the following commands, replace the example values with your actual values. The examples use `tsminst1` for the server instance user ID, `/tsminst1` for the server instance directory, and `/home/tsminst1` as the server instance users home directory.

1. Set the IBM Spectrum Protect API environment-variable configuration for the database instance:
   a. Log in by using the `tsminst1` user ID.
   b. When user `tsminst1` is logged in, ensure that the DB2 environment is properly initialized. The DB2 environment is initialized by running the `/home/tsminst1/sqllib/db2profile` script, which normally runs automatically from the profile of the user ID. Ensure the `.profile` file exists in the instance users home directory, for example, `/home/tsminst1/.profile`. If `.profile` does not run the `db2profile` script, add the following lines:

```
if [ -f /home/tsminst1/sqllib/db2profile ]; then
     . /home/tsminst1/sqllib/db2profile
  fi
```

   c. In the `instance_directory/sqllib/userprofile` file, add the following lines:

```
DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
DSMI_DIR=server_bin_directory/dbbkapi
DSMI_LOG=server_instance_directory
export DSMI_CONFIG DSMI_DIR DSMI_LOG
```

   where:
   - *instance_directory* is the home directory of the server instance user.
   - *server_instance_directory* is the server instance directory.
   - *server_bin_directory* is the server bin directory. The default location is `/opt/tivoli/tsm/server/bin`.

   In the `instance_directory/sqllib/usercshrc` file, add the following lines:

```
setenv DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
setenv DSMI_DIR=server_bin_directory/dbbkapi
setenv DSMI_LOG=server_instance_directory
```

2. Log off and log in again as `tsminst1`, or issue this command:

```
. ~/.profile
```

Tip: Ensure that you enter a space after the initial dot (.) character.

3. Create a file that is named `tsmdbmgr.opt` in the *server_instance* directory, which is in the `/tsminst1` directory in this example, and add the following line:

   ```
   SERVERNAME TSMDBMGR_TSMINST1
   ```

   Remember: The value for `SERVERNAME` must be consistent in the `tsmdbmgr.opt` and `dsm.sys` files.

4. As root user, add the following lines to the IBM Spectrum Protect API `dsm.sys` configuration file. By default, the `dsm.sys` configuration file is in the following default location:

   server_bin_directory`/dbbkapi/dsm.sys`

   ```
   servername TSMDBMGR_TSMINST1
   commmethod tcpip
   tcpserveraddr localhost
   tcpport 1500
   errorlogname /tsminst1/tsmdbmgr.log
   nodename $$_TSMDBMGR_$$
   ```

   where

   - *servername* matches the `servername` value in the `tsmdbmgr.opt` file.
   - *commethod* specifies the client API that is used to contact the server for database backup. This value can be `tcpip` or `sharedmem`. For more information about shared memory, see step 5.
   - *tcpserveraddr* specifies the server address that the client API uses to contact the server for database backup. To ensure that the database can be backed up, this value must be `localhost`.
   - *tcpport* specifies the port number that the client API uses to contact the server for database backup. Ensure that you enter the same `tcpport` value that is specified in the `dsmserv.opt` server options file.
   - *errorlogname* specifies the error log where the client API logs errors that are encountered during a database backup. This log is typically in the server instance directory. However, this log can be placed in any location where the instance user ID has write-permission.
   - *nodename* specifies the node name that the client API uses to connect to the server during a database backup. To ensure that the database can be backed up, this value must be `$$_TSMDBMGR_$$`.

   Attention: Do not add the `PASSWORDACCESS generate` option to the `dsm.sys` configuration file. This option can cause the database backup to fail.

5. Optional: Configure the server to back up the database by using shared memory. In this way, you might be able to reduce the processor load and improve throughput. Complete the following steps:
   a. Review the `dsmserv.opt` file. If the following lines are not in the file, add them:

      ```
      commmethod sharedmem
      shmport port_number
      ```

      where *port_number* specifies the port to be used for shared memory.

   b. In the `dsm.sys` configuration file, locate the following lines:

      ```
      commmethod tcpip
      tcpserveraddr localhost
      tcpport port_number
      ```

Replace the specified lines with the following lines:

```
commmethod sharedmem
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

Configuring server options for server database maintenance

To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.

About this task

Table and index reorganization requires significant processor resources, active log space, and archive log space. Because database backup takes precedence over reorganization, select the time and duration for reorganization to ensure that the processes do not overlap and reorganization can complete.

You can optimize index and table reorganization for the server database. In this way, you can help to avoid unexpected database growth and performance issues. For instructions, see technote 1683633.

If you update these server options while the server is running, you must stop and restart the server before the updated values take effect.

Procedure

1. Modify the server options.

   Edit the server options file, `dsmserv.opt`, in the server instance directory. Follow these guidelines when you edit the server options file:
   - To enable an option, remove the asterisk at the beginning of the line.
   - Enter an option on any line.
   - Enter only one option per line. The entire option with its value must be on one line.
   - If you have multiple entries for an option in the file, the server uses the last entry.

   To view available server options, see the sample file, `dsmserv.opt.smp`, in the `/opt/tivoli/tsm/server/bin` directory.
2. If you plan to use data deduplication, enable the **ALLOWREORGINDEX** server option. Add the following option and value to the server options file:

   ```
   allowreorgindex yes
   ```
3. Set the **REORGBEGINTIME** and **REORGDURATION** server options to control when reorganization starts and how long it runs. Select a time and duration so that reorganization runs when you expect that the server is least busy. These server options control both table and index reorganization processes.
   1. Set the time for reorganization to start by using the **REORGBEGINTIME** server option. Specify the time by using the 24-hour system. For example, to set the start time for reorganization as 8:30 p.m., specify the following option and value in the server options file:

      ```
      reorgbegintime 20:30
      ```

b. Set the interval during which the server can start reorganization. For example, to specify that the server can start reorganization for four hours after the time set by the **REORGBEGINTIME** server option, specify the following option and value in the server options file:

```
reorgduration 4
```

4. If the server was running while you updated the server options file, stop and restart the server.

Related information:

- ALLOWREORGINDEX
- ALLOWREORGTABLE
- REORGBEGINTIME
- REORGDURATION

Starting the server instance

You can start the server by using the instance user ID, which is the preferred method, or the root user ID.

Before you begin

Ensure that you set access permissions and user limits correctly. For instructions, see "Verifying access rights and user limits" on page 74.

About this task

When you start the server by using the instance user ID, you simplify the setup process and avoid potential issues. However, in some cases, it might be necessary to start the server with the root user ID. For example, you might want to use the root user ID to ensure that the server can access specific devices. You can set up the server to start automatically by using either the instance user ID or the root user ID.

If you must complete maintenance or reconfiguration tasks, start the server in maintenance mode.

Procedure

To start the server, take one of the following actions:

- Start the server by using the instance user ID.

  For instructions, see "Starting the server from the instance user ID" on page 75.
- Start the server by using the root user ID.

  For instructions about authorizing root user IDs to start the server, see Authorizing root user IDs to start the server (V7.1.1). For instructions about starting the server by using the root user ID, see Starting the server from the root user ID (V7.1.1).
- Start the server automatically.

  For instructions, see "Automatically starting servers on Linux systems" on page 76.
- Start the server in maintenance mode.

  For instructions, see "Starting the server in maintenance mode" on page 77.

Verifying access rights and user limits

Before you start the server, verify access rights and user limits.

About this task

If you do not verify user limits, also known as *ulimits*, you might experience server instability or a failure of the server to respond. You must also verify the system-wide limit for the maximum number of open files. The system-wide limit must be greater than or equal to the user limit.

Procedure

1. Verify that the server instance user ID has permissions to start the server.
2. For the server instance that you plan to start, ensure that you have authority to read and write files in the server instance directory. Verify that the `dsmserv.opt` file exists in the server instance directory, and that the file includes parameters for the server instance.
3. If the server is attached to a tape drive, medium changer, or removable media device, and you plan to start the server by using the instance user ID, grant read/write access to the instance user ID for these devices. To set permissions, take one of the following actions:
   - If the system is dedicated to IBM Spectrum Protect and only the IBM Spectrum Protect administrator has access, make the device special file world-writable. On the operating system command line, issue the following command:

     ```
     chmod +w /dev/rmtX
     ```
   - If the system has multiple users, you can restrict access by making the IBM Spectrum Protect instance user ID the owner of the special device files. On the operating system command line, issue the following command:

     ```
     chmod u+w /dev/rmtX
     ```
   - If multiple user instances are running on the same system, change the group name, for example TAPEUSERS, and add each IBM Spectrum Protect instance user ID to that group. Then, change the ownership of the device special files to belong to the group TAPEUSERS, and make them group-writable. On the operating system command line, issue the following command:

     ```
     chmod g+w /dev/rmtX
     ```
4. If you are using the IBM Spectrum Protect device driver and the **`autoconf`** utility, use the **`-a`** option to grant read/write access to the instance user ID.
5. To prevent server failures during interaction with DB2, tune the kernel parameters.

   For instructions about tuning kernel parameters, see "Tuning kernel parameters" on page 59.
6. Verify the following user limits based on the guidelines in the table.

Table 17. User limit (ulimit) values

User limit type	Preferred value	Command to query value
Maximum size of core files created	Unlimited	`ulimit -Hc`
Maximum size of a data segment for a process	Unlimited	`ulimit -Hd`
Maximum file size	Unlimited	`ulimit -Hf`

Table 17. User limit (ulimit) values (continued)

User limit type	Preferred value	Command to query value
Maximum number of open files	65536	`ulimit -Hn`
Maximum amount of processor time in seconds	Unlimited	`ulimit -Ht`

To modify user limits, follow the instructions in the documentation for your operating system.

Tip: If you plan to start the server automatically by using a script, you can set the user limits in the script.

7. Ensure that the user limit of maximum user processes (the `nproc` setting) is set to the minimum suggested value of 16384.
   a. To verify the current user limit, issue the `ulimit -Hu` command by using the instance user ID. For example:
   ```
   [user@Machine ~]$ ulimit -Hu
   16384
   ```
   b. If the limit of maximum user processes is not set to 16384, set the value to 16384.

      Add the following line to the `/etc/security/limits.conf` file:
      ```
      instance_user_id       -    nproc          16384
      ```
      where *instance_user_id* specifies the server instance user ID.

      If the server is installed on the Red Hat Enterprise Linux 6 operating system, set the user limit by editing the `/etc/security/limits.d/90-nproc.conf` file in the `/etc/security/limits.d` directory. This file overrides the settings in the `/etc/security/limits.conf` file.

      Tip: The default value for the user limit of maximum user processes has changed on some distributions and versions of the Linux operating system. The default value is 1024. If you do not change the value to the minimum suggested value of 16384, the server might fail or hang.

Starting the server from the instance user ID

To start the server from the instance user ID, log in with the instance user ID and issue the appropriate command from the server instance directory.

Before you begin

Ensure that access rights and user limits are set correctly. For instructions, see “Verifying access rights and user limits” on page 74.

Procedure

1. Log in to the system where IBM Spectrum Protect is installed by using the instance user ID for the server.
2. If you do not have a user profile that runs the `db2profile` script, issue the following command:
   ```
   . /home/tsminst1/sqllib/db2profile
   ```
   Tip: For instructions about updating the user ID login script to run the `db2profile` script automatically, see the DB2 documentation.

3. Start the server by issuing the following command on one line from the server instance directory:

   ```
   usr/bin/dsmserv
   ```

   Tip: The command runs in the foreground so that you can set an administrator ID and connect to the server instance.

   For example, if the name of the server instance is `tsminst1` and the server instance directory is `/tsminst1`, you can start the instance by issuing the following commands:

   ```
   cd /tsminst1
    . ~/sqllib/db2profile
   /usr/bin/dsmserv
   ```

Automatically starting servers on Linux systems

To automatically start a server on a Linux operating system, use the **`dsmserv.rc`** script.

Before you begin

Ensure that kernel parameters are set correctly. For instructions, see "Tuning kernel parameters" on page 59.

Ensure that the server instance runs under the instance owner user ID.

Ensure that access rights and user limits are set correctly. For instructions, see "Verifying access rights and user limits" on page 74.

About this task

The **`dsmserv.rc`** script is in the server installation directory, for example, `/opt/tivoli/tsm/server/bin`.

The **`dsmserv.rc`** script can be used either to start the server manually or to start the server automatically by adding entries to the `/etc/rc.d/init.d` directory. The script works with Linux utilities such as **`CHKCONFIG`** and **`SERVICE`**.

Procedure

For each server instance that you want to automatically start, complete the following steps:

1. Place a copy of the **`dsmserv.rc`** script in the `/init.d` directory, for example, `/etc/rc.d/init.d`.

   Ensure that you change only the copy of the script. Do not change the original script.
2. Rename the script copy so that it matches the name of the server instance owner, for example, `tsminst1`.

   The script was created under the assumption that the server instance directory is *`home_directory`*`/tsminst1`, for example: `/home/tsminst1/tsminst1`.
3. If the server instance directory is not *`home_directory`*`/tsminst1`, locate the following line in the script copy:

   ```
   instance_dir="${instance_home}/tsminst1"
   ```

   Change the line so that it points to your server instance directory, for example:

   ```
   instance_dir="/tsminst1"
   ```

4. In the script copy, locate the following line:

   ```
   # pidfile: /var/run/dsmserv_instancename.pid
   ```

   Change the instance name value to the name of the server instance owner. For example, if the server instance owner is tsminst1, update the line as shown:

   ```
   # pidfile: /var/run/dsmserv_tsminst1.pid
   ```

5. Configure the run level in which the server automatically starts. By using tools such as the **CHKCONFIG** utility, specify a value that corresponds to a multiuser mode, with networking turned on. Typically, the run level to use is 3 or 5, depending on the operating system and its configuration. For more information about multiuser mode and run levels, see the documentation for your operating system.
6. To start or stop the server, issue one of the following commands:
   - To start the server:

     ```
     service tsminst1 start
     ```

   - To stop the server:

     ```
     service tsminst1 stop
     ```

Example

This example uses the following values:
- The instance owner is tsminst1.
- The server instance directory is `/home/tsminst1/tsminst1`.
- The **dsmserv.rc** script copy is named `tsminst1`.
- The **CHKCONFIG** utility is used to configure the script to start at run levels 3, 4, and 5.

```
cp /opt/tivoli/tsm/server/bin/dsmserv.rc /etc/rc.d/init.d/tsminst1
sed -i 's/dsmserv_instancename.pid/dsmserv_tsminst1.pid/' /etc/rc.d/init.d/tsminst1
chkconfig --list tsminst1
service tsminst1 supports chkconfig, but is not referenced in
any runlevel (run 'chkconfig --add tsminst1')
chkconfig --add tsminst1
chkconfig --list tsminst1
tsminst1 0:off 1:off 2:off 3:off 4:off 5:off 6:off
chkconfig --level 345 tsminst1 on
chkconfig --list tsminst1
tsminst1 0:off 1:off 2:off 3:on 4:on 5:on 6:off
```

Related reference:

Server startup script: dsmserv.rc

Starting the server in maintenance mode

You can start the server in maintenance mode to avoid disruptions during maintenance and reconfiguration tasks.

About this task

Start the server in maintenance mode by running the **DSMSERV** utility with the **MAINTENANCE** parameter.

The following operations are disabled in maintenance mode:
- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server

- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:

- You do not have to edit the server options file, `dsmserv.opt`, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserv maintenance
```

Tip: To view a video about starting the server in maintenance mode, see Starting a server in maintenance mode.

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the **HALT** command:
   ```
   halt
   ```
2. Start the server by using the method that you use in production mode.

Operations that were disabled during maintenance mode are reenabled.

Stopping the server

You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.

About this task

To stop the server, issue the following command from the IBM Spectrum Protect command line:

```
halt
```

If you cannot connect to the server with an administrative client and you want to stop the server, you must cancel the process by using the **kill** command with the process ID number (pid). The pid is displayed at initialization.

Important: Before you issue the **kill** command, ensure that you know the correct process ID for the IBM Spectrum Protect server.
The `dsmserv.v6lock` file, in the directory from which the server is running, can be used to identify the process ID of the process to kill. To display the file, enter:

```
cat /instance_dir/dsmserv.v6lock
```

Issue the following command to stop the server:

```
kill -23 dsmserv_pid
```

where *dsmserv_pid* is the process ID number.

Registering licenses

Immediately register any IBM Spectrum Protect licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.

About this task

Use the **REGISTER LICENSE** command for this task.

Example: Register a license

Register the base IBM Spectrum Protect license.

```
register license file=tsmbasic.lic
```

Specifying a device class in preparation for database backups

To prepare the system for automatic and manual database backups, you must specify the device class to be used.

Before you begin

Ensure that you have defined a tape or file device class.

About this task

Complete the following steps to set up your system for database backups.

Procedure

1. If you did not use the configuration wizard (`dsmicfgx`) to configure the server, ensure that you have completed the steps to manually configure the system for database backups.
2. Select the device class to be used for backups of the database. Issue the following command from an IBM Spectrum Protect administrative command line.

   ```
   set dbrecovery device_class_name
   ```

   The device class that you specify is used by the database manager for database backups. If you do not specify a device class with the **SET DBRECOVERY** command, the backup fails.

Example

For example, to specify that the **DBBACK** device class is to be used, issue this command:

```
set dbrecovery dbback
```

Running multiple server instances on a single system

You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.

Multiply the memory and other system requirements for one server by the number of instances planned for the system.

The set of files for one instance of the server is stored separately from the files used by another server instance on the same system. Use the steps in “Creating the server instance” on page 64 for each new instance, including creation of the new instance user.

To manage the system memory that is used by each server, use the `DBMEMPERCENT` server option to limit the percentage of system memory. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.

You can upgrade directly from either V6.3 to V7.1. See the upgrade section (Chapter 5, “Upgrading to V8.1,” on page 87) for more details. When you upgrade and have multiple servers on your system, you must run the installation wizard only once. The installation wizard collects the database and variables information for all of your original server instances.

If you upgrade from IBM Spectrum Protect V6.3 to V8.1.2 and have multiple servers on your system, all instances that exist in DB2 V9.7 are dropped and recreated in DB2 V11.1. The wizard issues the `db2 upgrade` *`db dbname`* command for each database. The database environment variables for each instance on your system are also reconfigured during the upgrade process.

Related tasks:

➡ Running multiple server instances on a single system (V7.1.1)

Monitoring the server

When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Procedure

1. Monitor the active log to ensure that the size is correct for the workload that is handled by the server instance.

   When the server workload reaches its typical expected level, the space that is used by the active log is 80% - 90% of the space that is available to the active log directory. At that point, you might need to increase the amount of space. Whether you must increase the space depends on the types of transactions in the server workload. Transaction characteristics affect how the active log space is used.

   The following transaction characteristics can affect the space usage in the active log:

   - The number and size of files in backup operations

  - Clients such as file servers that back up large numbers of small files can cause large numbers of transactions that are completed quickly. The transactions might use a large amount of space in the active log, but for a short time.
  - Clients such as a mail server or a database server that back up large amounts of data in few transactions can cause small numbers of transactions that take a long time to complete. The transactions might use a small amount of space in the active log, but for a long time.
- Network connection types
  - Backup operations that occur over fast network connections cause transactions that complete more quickly. The transactions use space in the active log for a shorter time.
  - Backup operations that occur over relatively slower connections cause transactions that take a longer time to complete. The transactions use space in the active log for a longer time.

If the server is handling transactions with a wide variety of characteristics, the space that is used for the active log might increase and decrease significantly over time. For such a server, you might need to ensure that the active log typically has a smaller percentage of its space used. The extra space allows the active log to grow for transactions that take a long time to complete.

2. Monitor the archive log to ensure that space is always available.

   Remember: If the archive log becomes full, and the failover archive log becomes full, the active log can become full, and the server stops. The goal is to make enough space available to the archive log so that it never uses all its available space.
   You are likely to notice the following pattern:

   a. Initially, the archive log grows rapidly as typical client-backup operations occur.
   b. Database backups occur regularly, either as scheduled or done manually.
   c. After at least two full database backups occur, log pruning occurs automatically. The space that is used by the archive log decreases when the pruning occurs.
   d. Normal client operations continue, and the archive log grows again.
   e. Database backups occur regularly, and log pruning occurs as often as full database backups occur.

   With this pattern, the archive log grows initially, decreases, and then might grow again. Over time, as normal operations continue, the amount of space that is used by the archive log should reach a relatively constant level.

   If the archive log continues to grow, consider taking one or both of these actions:

   - Add space to the archive log. You might need to move the archive log to a different file system.
   - Increase the frequency of full database backups, so that log pruning occurs more frequently.

3. If you defined a directory for the failover archive log, determine whether any logs get stored in that directory during normal operations. If the failover log space is being used, consider increasing the size of the archive log. The goal is that the failover archive log is used only under unusual conditions, not in normal operation.

Chapter 4. Installing an IBM Spectrum Protect server fix pack

IBM Spectrum Protect maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.

Before you begin

To install a fix pack or interim fix to the server, install the server at the level on which you want to run it. You do not have to start the server installation at the base release level. For example, if you currently have V8.1.1 installed, you can go directly to the latest fix pack for V8.1. You do not have to start with the V8.1.0 installation if a maintenance update is available.

You must have the IBM Spectrum Protect license package installed. The license package is provided with the purchase of a base release. When you download a fix pack or interim fix from Fix Central, install the server license that is available on the Passport Advantage website. To display messages and help in a language other than US English, install the language package of your choice.

If you upgrade the server to V8.1.2 or later, and then revert the server to a level that is earlier than V8.1.2, you must restore the database to a point in time before the upgrade. During the upgrade process, complete the required steps to ensure that the database can be restored: back up the database, the volume history file, the device configuration file, and the server options file. For more information, see Chapter 6, “Reverting from Version 8.1.2 to a previous server,” on page 95.

If you are using the client management service, ensure that you upgrade it to the same version as the IBM Spectrum Protect server.

Ensure that you retain the installation media from the base release of the installed server. If you installed IBM Spectrum Protect from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Visit the IBM Support Portal for the following information:
- A list of the latest maintenance and download fixes. Click **Downloads** and apply any applicable fixes.
- Details about obtaining a base license package. Search for **Downloads > Passport Advantage**.
- Supported platforms and system requirements. Search for **IBM Spectrum Protect supported operating systems**.

Ensure that you upgrade the server before you upgrade backup-archive clients. If you do not upgrade the server first, communication between the server and clients might be interrupted.

Attention: Do not alter the DB2 software that is installed with IBM Spectrum Protect installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Procedure

To install a fix pack or interim fix, complete the following steps:

1. Back up the database. The preferred method is to use a snapshot backup. A snapshot backup is a full database backup that does not interrupt any scheduled database backups. For example, issue the following IBM Spectrum Protect administrative command:
   ```
   backup db type=dbsnapshot devclass=tapeclass
   ```
2. Back up the device configuration information. Issue the following IBM Spectrum Protect administrative command:
   ```
   backup devconfig filenames=file_name
   ```

   where *file_name* specifies the name of the file in which to store device configuration information.
3. Save the volume history file to another directory or rename the file. Issue the following IBM Spectrum Protect administrative command:
   ```
   backup volhistory filenames=file_name
   ```

   where *file_name* specifies the name of the file in which to store the volume history information.
4. Save a copy of the server options file, typically named `dsmserv.opt`. The file is in the server instance directory.
5. Halt the server before installing a fix pack or interim fix. Use the **`HALT`** command.
6. Ensure that extra space is available in the installation directory. The installation of this fix pack might require additional temporary disk space in the installation directory of the server. The amount of additional disk space can be as much as that required for installing a new database as part of an IBM Spectrum Protect installation. The IBM Spectrum Protect installation wizard displays the amount of space that is required for installing the fix pack and the available amount. If the required amount of space is greater than the available amount, the installation stops. If the installation stops, add the required disk space to the file system and restart the installation.
7. Log in as the root user.
8. Obtain the package file for the fix pack or interim fix that you want to install from the IBM Support Portal, Passport Advantage, or Fix Central.
9. Change to the directory where you placed the executable file and complete the following steps.

   Tip: The files are extracted to the current directory. Ensure that the executable file is in the directory where you want the extracted files to be located.
   a. Change file permissions by entering the following command:
      ```
      chmod a+x 8.x.x.x-IBM-SPSRV-platform.bin
      ```

      where *platform* denotes the architecture that IBM Spectrum Protect is to be installed on.
   b. Issue the following command to extract the installation files:
      ```
      ./8.x.x.x-IBM-SPSRV-platform.bin
      ```
10. Select one of the following ways of installing IBM Spectrum Protect.

Important: After a fix pack is installed, it is not necessary to go through the configuration again. You can stop after completing the installation, fix any errors, then restart your servers.

Install the IBM Spectrum Protect software by using one of the following methods:

Installation wizard

Follow the instructions for your operating system:

"Installing IBM Spectrum Protect by using the installation wizard" on page 54

Tip: After you start the wizard, in the IBM Installation Manager window, click the **Update** icon; do not click the **Install** or **Modify** icon.

Command line in console mode

Follow the instructions for your operating system:

"Installing IBM Spectrum Protect by using console mode" on page 55

Silent mode

Follow the instructions for your operating system:

"Installing IBM Spectrum Protect in silent mode" on page 56

Tip: If you have multiple server instances on your system, run the installation wizard only once. The installation wizard upgrades all server instances.

Results

Correct any errors that are detected during the installation process.

If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click **File** > **View Log**. To collect log files, from the IBM Installation Manager tool, click **Help** > **Export Data for Problem Analysis**.

If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

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

Installing an IBM Spectrum Protect fix pack

Chapter 5. Upgrading to V8.1

To take advantage of new product features and updates, upgrade the IBM Spectrum Protect server to Version 8.1.2.

Before you begin

Upgrade the IBM Spectrum Protect server before you update clients. For more information, see:

What you should know about security before you install or upgrade the server

About this task

To upgrade the server on the same operating system, see the upgrade instructions. For instructions about migrating the server to a different operating system, see IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions.

Table 18. Upgrade instructions

To upgrade from this version	To this version	See this information
V8.1	V8.1 fix pack or interim fix	Chapter 4, "Installing an IBM Spectrum Protect server fix pack," on page 83
V7.1	V8.1	"Installing the server and verifying the upgrade" on page 90
V7.1	V8.1 fix pack or interim fix	Chapter 4, "Installing an IBM Spectrum Protect server fix pack," on page 83
V5.5, V6.2, or V6.3	V8.1	IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions

An upgrade from V7 to V8.1 takes approximately 20 - 50 minutes. Your environment might produce different results from the results that were obtained in the labs.

For information about upgrades in a clustered environment, see "Upgrading the server in a clustered environment" on page 93.

To revert to an earlier version of the server after an upgrade or migration, you must have a full database backup and the installation software for the original server. You must also have the following key configuration files:

- Volume history file
- Device configuration file
- Server options file

Related information:

IBM Spectrum Protect Upgrade and Migration Process - Frequently Asked Questions

Upgrading to V8.1

You can upgrade the server directly from V7.1 to V8.1. You do not have to uninstall V7.1.

Before you begin

Ensure that you retain the installation media from the server base release that you are upgrading. If you installed the server components from a DVD, ensure that the DVD is available. If you installed the server components from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Tip: DVDs are no longer available with V8.1 and later.

Procedure

To upgrade the server to V8.1, complete the following tasks:
1. "Planning the upgrade"
2. "Preparing the system" on page 89
3. "Installing the server and verifying the upgrade" on page 90

Planning the upgrade

Before you upgrade the server from V7.1 to V8.1, you must review the relevant planning information, such as system requirements and release notes. Then, select an appropriate day and time to upgrade the system so that you can minimize the impact on production operations.

About this task

In lab tests, the process of upgrading the server from V7.1 to V8.1 took 14 - 45 minutes. The results that you achieve might differ, depending on your hardware and software environment, and the size of the server database.

Procedure

1. Review the hardware and software requirements:

   "Minimum system requirements" on page 21

   For the latest updates related to system requirements, see the IBM Spectrum Protect support website at technote 1243309.
2. For special instructions or specific information for your operating system, review the release notes (http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.2/srv.common/r_relnotes_srv.html) and readme files for server components.
3. Select an appropriate day and time to upgrade your system to minimize the impact on production operations. The amount of time that is required to update the system depends on the database size and many other factors. When you start the upgrade process, clients cannot connect to the server until the new software is installed and any required licenses are registered again.

4. If you are upgrading the server from V6 or V7 to V8.1, verify that you have the system ID and password for the DB2 instance of the IBM Spectrum Protect server. These credentials are required to upgrade the system.

Preparing the system

To prepare the system for the upgrade from V7.1 to V8.1, you must gather information about each DB2 instance. Then, back up the server database, save key configuration files, cancel sessions, and stop the server.

Procedure

1. Log on to the computer where the server is installed.

   Ensure that you are logged on with the instance user ID.
2. Obtain a list of DB2 instances. Issue the following system command:

   ```
   /opt/tivoli/tsm/db2/instance/db2ilist
   ```

   The output might be similar to the following example:

   ```
   tsminst1
   ```

   Ensure that each instance corresponds to a server that is running on the system.
3. For each DB2 instance, note the default database path, actual database path, database name, database alias, and any DB2 variables that are configured for the instance. Keep the record for future reference. This information is required to restore the V7.1 database.
4. Connect to the server by using an administrative user ID.
5. Back up the database by using the **BACKUP DB** command. The preferred method is to create a snapshot backup, which is a full database backup that does not interrupt scheduled database backups. For example, you can create a snapshot backup by issuing the following command:

   ```
   backup db type=dbsnapshot devclass=tapeclass
   ```
6. Back up the device configuration information to another directory by issuing the following administrative command:

   ```
   backup devconfig filenames=file_name
   ```

   where *file_name* specifies the name of the file in which to store device configuration information.

   Tip: If you decide to restore the V7.1 database, this file is required.
7. Back up the volume history file to another directory. Issue the following administrative command:

   ```
   backup volhistory filenames=file_name
   ```

   where *file_name* specifies the name of the file in which to store the volume history information.

   Tip: If you decide to restore the V7.1 database, this file is required.
8. Save a copy of the server options file, which is typically named `dsmserv.opt`. The file is in the server instance directory.
9. Prevent activity on the server by disabling new sessions. Issue the following administrative commands:

   ```
   disable sessions client
   disable sessions server
   ```

10. Verify whether any sessions exist, and notify the users that the server will be stopped. To check for existing sessions, issue the following administrative command:
    ```
    query session
    ```
11. Cancel sessions by issuing the following administrative command:
    ```
    cancel session all
    ```

    This command cancels all sessions except for your current session.
12. Stop the server by issuing the following administrative command:
    ```
    halt
    ```
13. Verify that the server is shut down and no processes are running.

    Issue the following command:
    ```
    ps -ef | grep dsmserv
    ```
14. In the server instance directory of your installation, locate the `NODELOCK` file and move it to another directory, where you are saving configuration files. The `NODELOCK` file contains the previous licensing information for your installation. This licensing information is replaced when the upgrade is complete.

Installing the server and verifying the upgrade

To complete the process of upgrading the server to V8.1, you must install the V8.1 server. Then, verify that the upgrade was successful by starting the server instance.

Before you begin

You must be logged on to the system by using the root user ID.

You can obtain the installation package from an IBM download site.

Set the system user limit for maximum file size to unlimited to ensure that the files can be downloaded correctly.

1. To query the maximum file size value, run the following command:
   ```
   ulimit -Hf
   ```
2. If the system user limit for maximum file size is not set to unlimited, change the setting to unlimited by completing the instructions in the documentation for your operating system.

About this task

By using the IBM Spectrum Protect installation software, you can install the following components:

- Server

  Tip: The database (DB2), the Global Security Kit (GSKit), and IBM Java Runtime Environment (JRE) are automatically installed when you select the server component.
- Server languages
- License
- Devices
- IBM Spectrum Protect for SAN
- Operations Center

Procedure

1. Download the appropriate package file from one of the following websites:
   - Download the server package from Passport Advantage or Fix Central.
   - For the most recent information, updates, and maintenance fixes, go to the IBM Support Portal.
2. Complete the following steps:
   1. Verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document for your product.
      - IBM Spectrum Protect technote 4042944
      - IBM Spectrum Protect Extended Edition technote 4042945
      - IBM Spectrum Protect for Data Retention technote 4042946
   2. Download the package file 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 to a directory that contains previously extracted files, or any other files.

      Also, ensure that you have executable permission for the package file.
   3. If necessary, run the following command to change the file permissions:

      ```
      chmod a+x package_name.bin
      ```

      where *package_name* is like the following example:

      ```
      8.1.x.000-IBM-SPSRV-Linuxs390x.bin
      8.1.x.000-IBM-SPSRV-Linuxx86_64.bin
      8.1.x.000-IBM-SPSRV-Linuxppc64le.bin
      ```

      In the examples, *8.1.x.000* represents the product release level.
   4. Extract the installation files by running the following command:

      ```
      ./package_name.bin
      ```

      The package is large. Therefore, the extraction takes some time.
3. Install the IBM Spectrum Protect software by using one of the following methods. Install the IBM Spectrum Protect license during the installation process.

   Tip: If you have multiple server instances on your system, install the IBM Spectrum Protect software only one time to upgrade all server instances.

   Installation wizard

   To install the server by using the graphical wizard of IBM Installation Manager, follow the instructions in "Installing IBM Spectrum Protect by using the installation wizard" on page 54.

   Ensure that your system meets the prerequisites for using the installation wizard. Then, complete the installation procedure. In the IBM Installation Manager window, click the **Update** or **Modify** icon.

   Installing the server by using the console mode

   To install the server by using the console mode, follow the instructions in "Installing IBM Spectrum Protect by using console mode" on page 55.

Review the information about installing the server in console mode and then complete the installation procedure.

Silent mode

To install the server by using silent mode, follow the instructions in "Installing IBM Spectrum Protect in silent mode" on page 56.

Review the information about installing the server in silent mode and then complete the installation procedure.

After you install the software, you do not have to reconfigure the system.

4. Correct any errors that are detected during the installation process.

   If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click **File** > **View Log**. To collect log files, from the IBM Installation Manager tool, click **Help** > **Export Data for Problem Analysis**.

   If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

   ```
   /var/ibm/InstallationManager/logs
   ```
5. Go to the IBM Support Portal to obtain fixes. Click **Fixes, updates, and drivers** and apply any applicable fixes.
6. Verify that the upgrade was successful:
   a. Start the server instance.

      For instructions, see "Starting the server instance" on page 73.

   b. Monitor the messages that the server issues as it starts. Watch for error and warning messages, and resolve any issues.

   c. Verify that you can connect to the server by using the administrative client. To start an administrative client session, run the following IBM Spectrum Protect administrative command:

      ```
      dsmadmc
      ```

   d. To obtain information about the upgraded system, run **QUERY** commands. For example, to obtain consolidated information about the system, run the following IBM Spectrum Protect administrative command:

      ```
      query system
      ```

      To obtain information about the database, run the following IBM Spectrum Protect administrative command:

      ```
      query db format=detailed
      ```
7. Register the licenses for the IBM Spectrum Protect server components that are installed on your system by running the **REGISTER LICENSE** command:

   ```
   register license file=installation_directory/server/bin/component_name.lic
   ```

   where *installation_directory* specifies the directory in which you installed the component, and *component_name* specifies the abbreviation for the component.

   For example, if you installed the server in the default directory, `/opt/tivoli/tsm`, run the following command to register the license:

   ```
   register license file=/opt/tivoli/tsm/server/bin/tsmbasic.lic
   ```

   For example, if you installed IBM Spectrum Protect Extended Edition in the `/opt/tivoli/tsm` directory, run the following command:

   ```
   register license file=/opt/tivoli/tsm/server/bin/tsmee.lic
   ```

For example, if you installed IBM Spectrum Protect for Data Retention in the `/opt/tivoli/tsm` directory, run the following command:

```
register license file=/opt/tivoli/tsm/server/bin/dataret.lic
```

Restriction:

You cannot use the IBM Spectrum Protect server to register licenses for the following products:

- IBM Spectrum Protect for Mail
- IBM Spectrum Protect for Databases
- IBM Spectrum Protect for ERP
- IBM Spectrum Protect for Space Management

The **REGISTER LICENSE** command does not apply to these licenses. The licensing for these products is done by IBM Spectrum Protect clients.

8. Optional: To install an extra language package, use the modify function of the IBM Installation Manager.
9. Optional: To upgrade to a newer version of a language package, use the update function of the IBM Installation Manager.

What to do next

You can authenticate passwords with the LDAP directory server, or authenticate passwords with the IBM Spectrum Protect server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

Upgrading the server in a clustered environment

To upgrade a server to V8.1.2 in a clustered environment, you must complete preparation and installation tasks. The procedures vary, depending on the operating system and release.

Procedure

Follow the procedure for your operating system, source release, and target release:

Table 19. Procedures for upgrading the server in a clustered environment on a Linux operating system

Source release	Target release	Procedure
V6 or V7	V8.1.2	Upgrading a server that is configured with Tivoli System Automation

Upgrading IBM Spectrum Protect to V8.1.2 in a clustered environment

To take advantage of new features in IBM Spectrum Protect, you can upgrade the IBM Spectrum Protect server that is installed on a Linux operating system in a clustered environment.

Procedure

To upgrade, follow the instructions in the configuring a Linux environment for clustering section.

Chapter 6. Reverting from Version 8.1.2 to a previous server

If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the IBM Spectrum Protect server with minimal loss of data.

Before you begin

You must have the following items from the earlier version of the server:

- Server database backup
- Volume history file
- Device configuration file
- Server options file

About this task

Use the same instructions whether you are reverting within releases or to an earlier release, for example, from 8.1.2 to 8.1.1 or from 8.1.2 to 7.1.2. The older version must match the version that you used before the upgrade to V8.1.

Attention: Specify the **REUSEDELAY** parameter to help prevent backup-archive client data loss when you revert the server to a previous version.

Steps for reverting to the previous server version

About this task

Complete the following steps on the system that has the V8.1 server.

Procedure

1. Halt the server to shut down all server operations by using the **HALT** command.
2. Remove the database from the database manager, then delete the database and recovery log directories.
   a. Manually remove the database. One way to remove it is by issuing this command:
   ```
   dsmserv removedb tsmdb1
   ```
   b. If you must reuse the space that is occupied by the database and recovery log directories, you can now delete these directories.
3. Use the uninstallation program to uninstall the V8.1 server. Uninstallation removes the server and the database manager, with their directories. For details, see Chapter 8, “Uninstalling IBM Spectrum Protect,” on page 103.
4. Stop the cluster service. Reinstall the version of the server program that you were using before the upgrade to V8.1.2. This version must match the version that your server was running when you created the database backup that you restore in a later step. For example, the server was at V7.1.7 before the

upgrade, and you intend to use the database backup that was in use on this server. You must install the V7.1.7 fix pack to be able to restore the database backup.

5. Configure the new server database by using the configuration wizard. To start the wizard, issue the following command:

   ```
   . /dsmicfgx
   ```

6. Ensure that no servers are running in the background.
7. Restore the database to a point in time before the upgrade.
8. Copy the following files to the instance directory.
   - Device configuration file
   - Volume history file
   - The server options file (typically `dsmserv.opt`)
9. If you enabled data deduplication for any FILE-type storage pools that existed before the upgrade, or if you moved data that existed before the upgrade into new storage pools while using the V8.1.2 server, you must complete additional recovery steps. For more details, see "Additional recovery steps if you created new storage pools or enabled data deduplication."
10. If the **REUSEDELAY** parameter setting on storage pools is less than the age of the database that you restored, restore volumes on any sequential-access storage pools that were reclaimed after that database backup. Use the **RESTORE VOLUME** command.

    If you do not have a backup of a storage pool, audit the reclaimed volumes by using the **AUDIT VOLUME** command, with the **FIX=YES** parameter to resolve inconsistencies. For example:

    ```
    audit volume volume_name fix=yes
    ```

11. If client backup or archive operations were completed using the V8.1 server, audit the storage pool volumes on which the data was stored.

Additional recovery steps if you created new storage pools or enabled data deduplication

If you created new storage pools, turned on data deduplication for any FILE-type storage pools, or did both while your server was running as a V8.1.2 server, you must complete more steps to return to the previous server version.

Before you begin

To complete this task, you must have a complete backup of the storage pool that was created before the upgrade to V8.1.2.

About this task

Use this information if you did either or both of the following actions while your server was running as a V8.1.2 server:

- You enabled the data deduplication function for any storage pools that existed before the upgrade to V8.1.2 program. Data deduplication applies only to storage pools that use a FILE device type.
- You created new primary storage pools after the upgrade *and* moved data that was stored in other storage pools into the new storage pools.

Complete these steps after the server is again restored to V7.

Procedure

- For each storage pool for which you enabled the data deduplication function, restore the entire storage pool by using the **RESTORE STGPOOL** command.
- For storage pools that you created after the upgrade, determine what action to take. Data that was moved from existing V8 storage pools into the new storage pools might be lost because the new storage pools no longer exist in your restored V8 server. Possible recovery depends on the type of storage pool:
  - If data was moved from V8 DISK-type storage pools into a new storage pool, space that was occupied by the data that was moved was probably reused. Therefore, you must restore the original V8 storage pools by using the storage pool backups that were created before the upgrade to V8.1.2.

    If *no* data was moved from V8 DISK-type storage pools into a new storage pool, then audit the storage pool volumes in these DISK-type storage pools.
  - If data was moved from V8 sequential-access storage pools into a new storage pool, that data might still exist and be usable in storage pool volumes on the restored V8 server. The data might be usable if the **REUSEDELAY** parameter for the storage pool was set to a value that prevented reclamation while the server was running as a V8.1.2 server. If any volumes were reclaimed while the server was running as a V8.1.2 server, restore those volumes from storage pool backups that were created before the upgrade to V8.1.2.

Reverting to a previous server version

Chapter 7. Reference: DB2 commands for IBM Spectrum Protect server databases

Use this list as reference when you are directed to issue DB2 commands by IBM support.

Purpose

After using the wizards to install and configure IBM Spectrum Protect, you seldom need to issue DB2 commands. A limited set of DB2 commands that you might use or be asked to issue are listed in Table 20. This list is supplemental material only and is not a comprehensive list. There is no implication that an IBM Spectrum Protect administrator will use it on a daily or ongoing basis. Samples of some commands are provided. Details of output are not listed.

For a full explanation of the commands described here and of their syntax, see the DB2 product information.

Table 20. DB2 commands

Command	Description	Example
`db2icrt`	Creates DB2 instances in the home directory of the instance owner. **Tip:** The IBM Spectrum Protect configuration wizard creates the instance used by the server and database. After a server is installed and configured through the configuration wizard, the **`db2icrt`** command is generally not used. This utility is in the `DB2DIR/instance` directory, where `DB2DIR` represents the installation location where the current version of the DB2 database system is installed.	Manually create an IBM Spectrum Protect instance. Enter the command on one line: `/opt/tivoli/tsm/db2/instance/ db2icrt -a server -u` *`instance_name instance_name`*
`db2set`	Displays DB2 variables.	List DB2 variables: `db2set`
`CATALOG DATABASE`	Stores database location information in the system database directory. The database can be located either on the local workstation or on a remote database partition server. The server configuration wizard takes care of any catalog needed for using the server database. Run this command manually, after a server is configured and running, only if something in the environment changes or is damaged.	Catalog the database: `db2 catalog database tsmdb1`
`CONNECT TO DATABASE`	Connects to a specified database for command-line interface (CLI) use.	Connect to the IBM Spectrum Protect database from a DB2 CLI: `db2 connect to tsmdb1`

Table 20. DB2 commands (continued)

Command	Description	Example

GET DATABASE CONFIGURATION	Returns the values of individual entries in a specific database configuration file.
	Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.
		Show the configuration information for a database alias:
		`db2 get db cfg for tsmdb1`
		Retrieve information in order to verify settings such as database configuration, log mode, and maintenance.
		`db2 get db config for tsmdb1 show detail`

GET DATABASE MANAGER CONFIGURATION	Returns the values of individual entries in a specific database configuration file. **Important:** This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of IBM Spectrum Protect server commands or procedures.	Retrieve configuration information for the database manager: `db2 get dbm cfg`
GET HEALTH SNAPSHOT	Retrieves the health status information for the database manager and its databases. The information returned represents a snapshot of the health state at the time the command was issued. IBM Spectrum Protect monitors the state of the database using the health snapshot and other mechanisms that are provided by DB2. There might be cases where the health snapshot or other DB2 documentation indicates that an item or database resource might be in an alert state. Such a case indicates that action must be considered to remedy the situation. IBM Spectrum Protect monitors the condition and responds appropriately. Not all declared alerts by the DB2 database are acted on.	Receive a report on DB2 health monitor indicators: `db2 get health snapshot for database on tsmdb1`

GRANT (Database Authorities)	Grants authorities that apply to the entire database rather than privileges that apply to specific objects within the database.	Grant access to the user ID itmuser:
		`db2 GRANT CONNECT ON DATABASE TO USER itmuser`
		`db2 GRANT CREATETAB ON DATABASE TO USER itmuser`

Table 20. DB2 commands (continued)

Command	Description	Example
`RUNSTATS`	Updates statistics about the characteristics of a table and associated indexes or statistical views. These characteristics include number of records, number of pages, and average record length. To see a table, issue this utility after updating or reorganizing the table. A view must be enabled for optimization before its statistics can be used to optimize a query. A view that is enabled for optimization is known as a statistical view. Use the DB2 **`ALTER VIEW`** statement to enable a view for optimization. Issue the **`RUNSTATS`** utility when changes to underlying tables substantially affect the rows returned by the view. **Tip:** The server configures DB2 to run the **`RUNSTATS`** command as needed.	Update statistics on a single table. `db2 runstats on table SCHEMA_NAME.TABLE_NAME with distribution and sampled detailed indexes all`
`SET SCHEMA`	Changes the value of the **`CURRENT SCHEMA`** special register, in preparation for issuing SQL commands directly through the DB2 CLI. **Tip:** A special register is a storage area that is defined for an application process by the database manager. It is used to store information that can be referenced in SQL statements.	Set the schema for IBM Spectrum Protect: `db2 set schema tsmdb1`
`START DATABASE MANAGER`	Starts the current database manager instance background processes. The server starts and stops the instance and database whenever the server starts and halts. **Important:** Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Start the database manager: `db2start`
`STOP DATABASE MANAGER`	Stops the current database manager instance. Unless explicitly stopped, the database manager continues to be active. This command does not stop the database manager instance if any applications are connected to databases. If there are no database connections, but there are instance attachments, the command forces the instance attachments to stop first. Then, it stops the database manager. This command also deactivates any outstanding database activations before stopping the database manager. This command is not valid on a client. The server starts and stops the instance and database whenever the server starts and halts. **Important:** Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.	Stop the database manager: `db2 stop dbm`

Chapter 8. Uninstalling IBM Spectrum Protect

You can use the following procedures to uninstall IBM Spectrum Protect. Before you remove IBM Spectrum Protect, ensure that you do not lose your backup and archive data.

Before you begin

Complete the following steps before you uninstall IBM Spectrum Protect:
- Complete a full database backup.
- Save a copy of the volume history and device configuration files.
- Store the output volumes in a safe location.

About this task

You can uninstall IBM Spectrum Protect by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

"Uninstalling IBM Spectrum Protect by using a graphical wizard"

"Uninstalling IBM Spectrum Protect in console mode" on page 104

"Uninstalling IBM Spectrum Protect in silent mode" on page 104

What to do next

See Chapter 2, "Installing the server components," on page 53 for installation steps to reinstall the IBM Spectrum Protect components.

Uninstalling IBM Spectrum Protect by using a graphical wizard

You can uninstall IBM Spectrum Protect by using the IBM Installation Manager installation wizard.

Procedure

1. Start the Installation Manager.

   In the directory where the Installation Manager is installed, go to the `eclipse` subdirectory (for example, `/opt/IBM/InstallationManager/eclipse`), and issue the following command:

   ```
   ./IBMIM
   ```
2. Click **Uninstall**.
3. Select **IBM Spectrum Protect server**, and click **Next**.
4. Click **Uninstall**.
5. Click **Finish**.

Uninstalling IBM Spectrum Protect in console mode

To uninstall IBM Spectrum Protect by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
   ```
   eclipse/tools
   ```
   For example:
   ```
   /opt/IBM/InstallationManager/eclipse/tools
   ```
2. From the `tools` directory, issue the following command:
   ```
   ./imcl -c
   ```
3. To uninstall, enter `5`.
4. Choose to uninstall from the IBM Spectrum Protect package group.
5. Enter `N` for Next.
6. Choose to uninstall the IBM Spectrum Protect server package.
7. Enter `N` for Next.
8. Enter `U` for Uninstall.
9. Enter `F` for Finish.

Uninstalling IBM Spectrum Protect in silent mode

To uninstall IBM Spectrum Protect in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the IBM Spectrum Protect server components. IBM Spectrum Protect includes a sample response file, `uninstall_response_sample.xml`, in the `input` directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

If you want to uninstall all IBM Spectrum Protect components, leave `modify="false"` set for each component in the response file. If you do not want to uninstall a component, set the value to `modify="true"`.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
   ```
   eclipse/tools
   ```
   For example:
   ```
   /opt/IBM/InstallationManager/eclipse/tools
   ```
2. From the `tools` directory, issue the following command, where *response_file* represents the response file path, including the file name:

```
./imcl -input response_file -silent
```

The following command is an example:

```
./imcl -input /tmp/input/uninstall_response.xml -silent
```

Uninstalling and reinstalling IBM Spectrum Protect

If you plan to manually reinstall IBM Spectrum Protect instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.

About this task

To manually uninstall and reinstall IBM Spectrum Protect, complete the following steps:

1. Make a list of your current server instances before proceeding to the uninstallation. Run the following command:

   ```
   /opt/tivoli/tsm/db2/instance/db2ilist
   ```

2. Run the following commands for every server instance:

   ```
   db2 attach to instance_name
   db2 get dbm cfg show detail
   db2 detach
   ```

   Keep a record of the database path for each instance.

3. Uninstall IBM Spectrum Protect. See Chapter 8, “Uninstalling IBM Spectrum Protect,” on page 103.
4. When you uninstall any supported version of IBM Spectrum Protect, including a fix pack, an instance file is created. The instance file is created to help reinstall IBM Spectrum Protect. Check this file and use the information when you are prompted for the instance credentials when reinstalling. In silent installation mode, you provide these credentials using the `INSTANCE_CRED` variable.

   You can find the instance file in the following location:

   ```
   /etc/tivoli/tsm/instanceList.obj
   ```

5. Reinstall IBM Spectrum Protect. See Chapter 2, “Installing the server components,” on page 53.

   If the `instanceList.obj` file does not exist, you need to recreate your server instances using the following steps:

   a. Recreate your server instances. See “Creating the server instance” on page 64.

      Tip: The installation wizard configures the server instances but you must verify that they exist. If they do not exist, you must manually configure them.

   b. Catalog the database. Log in to each server instance as the instance user, one at a time, and issue the following commands:

      ```
      db2 catalog database tsmdb1
      db2 attach to instance_name
      db2 update dbm cfg using dftdbpath instance_directory
      db2 detach
      ```

   c. Verify that the server instance was created successfully. Issue this command:

      ```
      /opt/tivoli/tsm/db2/instance/db2ilist
      ```

d. Verify that IBM Spectrum Protect recognizes the server instance by listing your directories. Your home directory appears if you did not change it. Your instance directory does appear if you used the configuration wizard. Issue this command:

```
db2 list database directory
```

If you see TSMDB1 listed, you can start the server.

Uninstalling IBM Installation Manager

You can uninstall IBM Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

Before you begin

Before you uninstall IBM Installation Manager, you must ensure that all packages that were installed by IBM Installation Manager are uninstalled. Close IBM Installation Manager before you start the uninstall process.

To view installed packages, issue the following command from a command line:

```
cd /opt/IBM/InstallationManager/eclipse/tools
./imcl listInstalledPackages
```

Procedure

To uninstall IBM Installation Manager, complete the following steps:

1. Open a command line and change directories to `/var/ibm/InstallationManager/uninstall`.
2. Issue the following command:

   ```
   ./uninstall
   ```

   Restriction: You must be logged in to the system as the `root` user ID.

Part 2. Installing and upgrading the Operations Center

The IBM Spectrum Protect Operations Center is the web-based interface for managing your storage environment.

Before you begin

Before you install and configure the Operations Center, review the following information:

- "System requirements for the Operations Center" on page 109
  - "Operations Center computer requirements" on page 110
  - "Hub and spoke server requirements" on page 110
  - "Operating system requirements" on page 113
  - "Web browser requirements" on page 114
  - "Language requirements" on page 114
  - "Requirements and limitations for IBM Spectrum Protect client management services" on page 115
- "Administrator IDs that the Operations Center requires" on page 117
- "IBM Installation Manager" on page 117
- "Installation checklist" on page 118
- "Obtaining the Operations Center installation package" on page 121

About this task

Table 21 lists the methods for installing or uninstalling the Operations Center and indicates where to find the associated instructions.

For information about upgrading the Operations Center, see Chapter 11, "Upgrading the Operations Center," on page 125.

Table 21. Methods for installing or uninstalling the Operations Center

Method	Instructions
Graphical wizard	• "Installing the Operations Center by using a graphical wizard" on page 122 • "Uninstalling the Operations Center by using a graphical wizard" on page 161
Console mode	• "Installing the Operations Center in console mode" on page 122 • "Uninstalling the Operations Center in console mode" on page 161
Silent mode	• "Installing the Operations Center in silent mode" on page 122 • "Uninstalling the Operations Center in silent mode" on page 162

Chapter 9. Planning to install the Operations Center

Before you install the Operations Center, you must understand the system requirements, the administrator IDs that the Operations Center requires, and the information that you must provide to the installation program.

About this task

From the Operations Center, you can manage the following primary aspects of the storage environment:

- IBM Spectrum Protect servers and clients
- Services such as backup and restore, archive and retrieve, and migrate and recall
- Storage pools and storage devices

The Operations Center includes the following features:

User interface for multiple servers

You can use the Operations Center to manage one or more IBM Spectrum Protect servers.

In an environment with multiple servers, you can designate one server as a *hub server* and the others as *spoke servers*. The hub server can receive alerts and status information from the spoke servers and present the information in a consolidated view in the Operations Center.

Alert monitoring

An *alert* is a notification of a relevant problem on the server and is triggered by a server message. You can define which server messages trigger alerts, and only those messages are reported as alerts in the Operations Center or in an email.

This alert monitoring can help you identify and track relevant problems on the server.

Convenient command-line interface

The Operations Center includes a command-line interface for advanced features and configuration.

System requirements for the Operations Center

Before you install the Operations Center, ensure that your system meets the minimum requirements.

Use the Operations Center System Requirements Calculator to estimate the system requirements for running the Operations Center and the hub and spoke servers that are monitored by the Operations Center.

Requirements that are verified during the installation

Table 22 on page 110 lists the prerequisite requirements that are verified during the installation and indicates where to find more information about these requirements.

Table 22. Requirements that are verified during the installation

Requirement	Details
Minimum memory requirement	"Operations Center computer requirements"
Operating system requirement	"Operating system requirements" on page 113
Host name for the computer where the Operations Center will be installed	"Installation checklist" on page 118
Requirements for the Operations Center installation directory	"Installation checklist" on page 118

Operations Center computer requirements

You can install the Operations Center on a computer that is also running IBM Spectrum Protect server or on a different computer. If you install the Operations Center on the same computer as a server, that computer must meet the system requirements for both the Operations Center and the server.

Resource requirements

The following resources are required to run the Operations Center:

- One processor core
- 4 GB of memory
- 1 GB of disk space

The hub and spoke servers that are monitored by the Operations Center require additional resources, as described in "Hub and spoke server requirements."

Hub and spoke server requirements

When you open the Operations Center for the first time, you must associate the Operations Center with one IBM Spectrum Protect server that is designated as the *hub server*. In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.

The spoke servers send alerts and status information to the hub server. The Operations Center shows you a consolidated view of alerts and status information for the hub server and any spoke servers.

If only one server is monitored by the Operations Center, that server is still called a hub server, even though no spoke servers are connected to it.

Table 23 indicates the version of IBM Spectrum Protect server that must be installed on the hub server and on each spoke server that is managed by the Operations Center.

Table 23. IBM Spectrum Protect server version requirements for hub and spoke servers

Operations Center	Version on the hub server	Version on each spoke server
V8.1.2	V8.1.2	V6.3.4 or later **Restriction:** Some Operations Center functions are not available for servers that use a version earlier than V8.1.2.

Number of spoke servers that a hub server can support

The number of spoke servers that a hub server can support depends on the configuration and on the version of IBM Spectrum Protect on each spoke server. However, a general guideline is that a hub server can support 10 - 20 V6.3.4 spoke servers but can support more V7.1 or later spoke servers.

Tips for designing the hub and spoke server configuration

In designing the hub and spoke configuration, especially consider the resource requirements for status monitoring. Also, consider how you want to group hub and spoke servers and whether you want to use multiple hub servers.

Use the Operations Center System Requirements Calculator to estimate the system requirements for running the Operations Center and the hub and spoke servers that are monitored by the Operations Center.

Primary factors that affect performance

The following factors have the most significant impact on the performance of the Operations Center:

- The processor and memory on the computer where the Operations Center is installed
- The system resources of the hub and spoke servers, including the disk system that is in use for the hub server database
- The number of client nodes and virtual machine file spaces that are managed by the hub and spoke servers
- The frequency at which data is refreshed in the Operations Center

How to group hub and spoke servers

Consider grouping hub and spoke servers by geographic location. For example, managing the servers within the same data center can help prevent issues that are caused by firewalls or by inadequate network bandwidth between different locations. If necessary, you can further divide servers according to one or more of the following characteristics:

- The administrator who manages the servers
- The organizational entity that funds the servers
- Server operating system
- The language in which the servers run

  Tip: If the hub and spoke servers are not running in the same language, you might see corrupted text in the Operations Center.

How to group hub and spoke servers in an enterprise configuration

In an enterprise configuration, a network of IBM Spectrum Protect servers are managed as a group. Changes that are made on the *configuration manager* can be distributed automatically to one or more *managed servers* in the network.

The Operations Center normally registers and maintains a dedicated administrator ID on the hub and spoke servers. This *monitoring administrator* must always have the same password on all the servers.

If you use an enterprise configuration, you can improve the process by which the administrator credentials are synchronized on spoke servers. To improve the performance and efficiency of maintaining the monitoring administrator ID, complete the following steps:

1. Designate the configuration manager server as the Operations Center hub server. During the hub server configuration, a monitoring administrator ID named IBM-OC-*hub_server_name* is registered.
2. On the hub server, add the monitoring administrator ID to a new or existing enterprise configuration profile. Issue the NOTIFY SUBSCRIBERS command to distribute the profile to the managed servers.
3. Add one or more of the managed servers as Operations Center spoke servers.

The Operations Center detects this configuration and allows the configuration manager to distribute and update the monitoring administrator ID on the spoke servers.

When to use multiple hub servers

If you have more than 10 - 20 V6.3.4 spoke servers, or if resource limitations require the environment to be partitioned, you can configure multiple hub servers, and connect a subset of the spoke servers to each hub server.

Restrictions:

- A single server cannot be both a hub server and a spoke server.
- Each spoke server can be assigned to only one hub server.
- Each hub server requires a separate instance of the Operations Center, each of which has a separate web address.

Tips for choosing a hub server

For the hub server, you must choose a server that has adequate resources and is located for minimal roundtrip network latency.

Attention: Do not use the same server as the hub server for multiple Operations Centers.

Use the following guidelines in deciding which server to designate as the hub server:

Choose a lightly loaded server

Consider a server that has a light load for operations such as client backup and archive. A lightly loaded server is also a good choice as the host system for the Operations Center.

Ensure that the server has the resources to handle both its typical server workload and the estimated workload for acting as the hub server.

Locate the server for minimal roundtrip network latency

Locate the hub server so that the network connection between the hub server and the spoke servers has a roundtrip latency that is no greater than 5 ms. This latency can typically be achieved when the servers are on the same local area network (LAN).

Networks that are poorly tuned, are heavily used by other applications, or have roundtrip latency much higher than 5 ms can degrade communications between the hub and spoke servers. For example, roundtrip latencies of 50 ms or higher can result in communication timeouts that cause spoke servers to disconnect or reconnect to the

Operations Center. Such high latencies might be experienced in long-distance, wide area network (WAN) communications.

If spoke servers are a long distance from the hub server and experience frequent disconnects in the Operations Center, you can increase the value of the `ADMINCOMMTIMEOUT` option on each server to reduce the problem.

Verify that the hub server meets the resource requirements for status monitoring

Status monitoring requires extra resources on each server on which it is enabled. The resources that are required depend primarily on the number of clients that are managed by the hub and spoke servers. Fewer resources are used on a hub server with a V7.1 or later spoke server than on a hub server with a V6.3.4 spoke server.

Verify that the hub server meets the resource requirements for processor usage, database space, archive log space, and I/O operations per second (IOPS) capacity.

A hub server with high IOPS capacity can handle a larger amount of incoming status data from spoke servers. Use of the following storage devices for the hub server database can help meet this capacity:

- An enterprise-level solid-state drive (SSD)
- An external SAN disk storage device with multiple volumes or multiple spindles under each volume

In an environment with fewer than 1000 clients, consider establishing a baseline capacity of 1000 IOPS for the hub server database if the hub server manages any spoke servers.

Determine whether your environment requires multiple hub servers

If more than 10,000 - 20,000 client nodes and virtual machine file spaces are managed by one set of hub and spoke servers, the resource requirements might exceed what the hub server has available, especially if the spoke servers are V6.3.4 servers. Consider designating a second server as a hub server and moving spoke servers to the new hub server to balance the load.

Operating system requirements

The Operations Center is available for AIX, Linux, and Windows systems.

You can run the Operations Center on the following systems:

- Linux on x86_64 systems:
  - Red Hat Enterprise Linux 6.7
  - Red Hat Enterprise Linux 7.1
  - SUSE Linux Enterprise Server 11, Service Pack 4 or later
  - SUSE Linux Enterprise Server 12
- Linux on System z (s390x 64-bit architecture) systems:
  - Red Hat Enterprise Linux 7.1
  - SUSE Linux Enterprise Server 12
- Linux on Power Systems (little endian) systems:
  - Red Hat Enterprise Linux 7 with the PPC64LE architecture

For the most up-to-date requirements information, see Software and Hardware Requirements.

Web browser requirements

The Operations Center can run in Apple, Google, Microsoft, and Mozilla web browsers.

For optimal viewing of the Operations Center in the web browser, ensure that the screen resolution for the system is set to a minimum of 1024 X 768 pixels.

For optimal performance, use a web browser that has good JavaScript performance, and enable browser caching.

The Operations Center can run in the following web browsers:

- Apple Safari on the iPad

  Restriction: If Apple Safari is running on iOS 8.x or iOS 9.x, you cannot use a self-signed certificate for secure communication with the Operations Center without extra configuration of the certificate. Use a certificate authority (CA) certificate, or configure the self-signed certificate as needed. For instructions, see Technote http://www.ibm.com/support/docview.wss?uid=swg21963153.
- Google Chrome 40 or later
- Microsoft Internet Explorer 11 or later
- Mozilla Firefox ESR 31 or later

To run the Operations Center in compliance with the National Institute of Standards and Technology (NIST) Special Publications (SP) 800-131A recommendation, communication between the Operations Center and the web browser must be secured by using the Transport Layer Security (TLS) 1.2 protocol. During installation, you specify whether SP 800-131A compliance is required and the level of compliance. If strict SP 800-131A compliance is specified during installation, the web browser must support TLS 1.2, and TLS 1.2 must be enabled.

The web browser displays an SSL error if strict SP 800-131A compliance is specified during installation, and the web browser does not meet the preceding requirements.

Language requirements

By default, the Operations Center uses the language that the web browser uses. However, the installation process uses the language that the operating system uses. Verify that the web browser and the operating system are set to the language that you require.

Table 24. Operations Center language values that you can use on Linux systems

Language	**Language option value**
Chinese, Simplified	zh_CN
Chinese, Simplified (GBK)	zh_CN.gb18030
Chinese, Simplified (UTF-8)	zh_CN.utf8
Chinese, Traditional (Big5)	Zh_TW
Chinese, Traditional (euc_tw)	zh_TW
Chinese, Traditional (UTF-8)	zh_TW.utf8
English, United States	en_US
English (UTF-8)	en_US.utf8
French	fr_FR

Table 24. Operations Center language values that you can use on Linux systems (continued)

Language	Language option value
French (UTF-8)	fr_FR.utf8
German	de_DE
German (UTF-8)	de_DE.utf8
Italian	it_IT
Italian (UTF-8)	it_IT.utf8
Japanese (EUC)	ja_JP
Japanese (UTF-8)	ja_JP.utf8
Korean	ko_KR
Korean (UTF-8)	ko_KR.utf8
Portuguese, Brazilian	pt_BR
Portuguese, Brazilian (UTF-8)	pt_BR.utf8
Russian	ru_RU
Russian (UTF-8)	ru_RU.utf8
Spanish	es_ES
Spanish (UTF-8)	es_ES.utf8

Requirements and limitations for IBM Spectrum Protect client management services

IBM Spectrum Protect client management services is a component that you install on backup-archive clients to collect diagnostic information such as client log files. Before you install the client management service on your system, you must understand the requirements and limitations.

In the documentation for the client management service, *client system* is the system where the backup-archive client is installed.

Diagnostic information can be collected only from Linux and Windows clients, but administrators can view the diagnostic information in the Operations Center on AIX, Linux, or Windows operating systems.

Requirements for the client management service

Verify the following requirements before you install the client management service:

- To remotely access the client, the Operations Center administrator must have system authority or one of the following client authority levels:
  - Policy authority
  - Client owner authority
  - Client node access authority
- Ensure that the client system meets the following requirements:
  - The client management service can be installed only on client systems that run on Linux or Windows operating systems:
    - Linux x86 64-bit operating systems that are supported for the backup-archive client.

    - Windows 32-bit and 64-bit operating systems that are supported for the backup-archive client.
  - Transport Layer Security (TLS) 1.2 must be installed for transmission of data between the client management service and Operations Center. Basic authentication is provided and data and authentication information are encrypted through the SSL channel. TLS 1.2 is automatically installed along with the necessary SSL certificates when you install the client management service.
- On Linux client systems, you must have root user authority to install the client management service.
- For client systems that can have multiple client nodes, such as Linux client systems, ensure that each node name is unique on the client system.

  Tip: After you install the client management service, you do not have to install it again because the service can discover multiple client options files.

Limitations of the client management service

The client management service provides basic services for collecting diagnostic information from backup-archive clients. The following limitations exist for the client management service:

- You can install the client management service only on systems with backup-archive clients, including backup-archive clients that are installed on data mover nodes for IBM Spectrum Protect for Virtual Environments: Data Protection for VMware.
- You cannot install the client management service on other IBM Spectrum Protect client components or products that do not have backup-archive clients.
- If the backup-archive clients are protected by a firewall, ensure that the Operations Center can connect to the backup-archive clients through the firewall by using the configured port for the client management service. The default port is 9028, but it can be changed.
- The client management service scans all client log files to locate entries for the previous 72-hour period.
- The Diagnosis page in the Operations Center provides basic troubleshooting information for backup-archive clients. However, for some backup issues, you might have to access the client system and obtain further diagnostic information.
- If the combined size of the client error log files and schedule log files on a client system is more than 500 MB, delays can occur in sending log records to the Operations Center. You can control the size of the log files by enabling log file pruning or wrapping by specifying the **`errorlogretention`** or **`errorlogmax`** client option.
- If you use the same client node name to connect to multiple IBM Spectrum Protect servers that are installed on the same server, you can view log files for only one of the client nodes.

For updates about the client management service, including requirements, limitations, and documentation updates, see technote 1963610.

Related tasks:

"Collecting diagnostic information with IBM Spectrum Protect client management services" on page 139

Administrator IDs that the Operations Center requires

An administrator must have a valid ID and password on the hub server to log in to the Operations Center. An administrator ID is also assigned to the Operations Center so that the Operations Center can monitor servers.

The Operations Center requires the following IBM Spectrum Protect administrator IDs:

Administrator IDs that are registered on the hub server

Any administrator ID that is registered on the hub server can be used to log in to the Operations Center. The authority level of the ID determines which tasks can be completed. You can create new administrator IDs by using the **`REGISTER ADMIN`** command.

Restriction: To use an administrator ID in a multiple-server configuration, the ID must be registered on the hub and spoke servers with the same password and authority level.

To manage authentication for these servers, consider using one of the following methods:

- A Lightweight Directory Access Protocol (LDAP) server
- The enterprise configuration functions to automatically distribute changes to the administrator definitions.

Monitoring administrator ID

When you initially configure the hub server, an administrator ID named IBM-OC-*server_name* is registered with system authority on the hub server and is associated with the initial password that you specify. This ID, which is sometimes called the *monitoring administrator*, is intended for use only by the Operations Center.

Do not delete, lock, or modify this ID. The same administrator ID with the same password is registered on the spoke servers that you add. The password is automatically changed on the hub and spoke servers every 90 days. You do not need to use or manage this password.

Restriction: The Operations Center maintains the monitoring administrator ID and password on spoke servers unless you use an enterprise configuration to manage these credentials. For more information about using an enterprise configuration to manage the credentials, see "Tips for designing the hub and spoke server configuration" on page 111.

IBM Installation Manager

The Operations Center uses IBM Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install the Operations Center. It must remain installed on the system so that the Operations Center can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering
An installable unit of a software product.

The Operations Center offering contains all of the media that IBM Installation Manager requires to install the Operations Center.

Package
The group of software components that are required to install an offering.

The Operations Center package contains the following components:

- IBM Installation Manager installation program
- Operations Center offering

Package group
A set of packages that share a common parent directory.

Repository
A remote or local storage area for data and other application resources.

The Operations Center package is stored in a repository on IBM Fix Central.

Shared resources directory
A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of the Operations Center.

Installation checklist

Before you install the Operations Center, you must verify certain information, such as the installation credentials, and you must determine the input to provide to IBM Installation Manager for the installation.

The following checklist highlights the information that you must verify or determine before you install the Operations Center, and Table 25 on page 119 describes the details of this information:

__ Verify the host name for the computer where the Operations Center is to be installed.
__ Verify the installation credentials.
__ Determine the Operations Center installation directory, if you do not want to accept the default path.
__ Determine the IBM Installation Manager installation directory, if you do not want to accept the default path.
__ Determine the port number to be used by the Operations Center web server, if you do not want to accept the default port number.
__ Determine the password for secure communications.
__ Determine whether secure communications must comply with the National Institute of Standards and Technology (NIST) Special Publications (SP) 800-131A recommendation.

Table 25. Information to verify or determine before you install the Operations Center

Information	**Details**
Host name for the computer where the Operations Center is to be installed.	The host name must meet the following criteria: • It must not contain double-byte character set (DBCS) characters or the underscore character (_). • Although the host name can contain the hyphen character (-), it cannot have a hyphen as the last character in the name.
Installation credentials	To install the Operations Center, you must use the following user account: • The root user

Operations Center installation directory	The Operations Center is installed in the `ui` subdirectory of the installation directory.
	The following path is the default path for the Operations Center installation directory:
	• `/opt/tivoli/tsm`
	For example, if you use this default path, the Operations Center is installed in the following directory:
	`/opt/tivoli/tsm/ui`
	The installation directory path must meet the following criteria:
	• The path must contain no more than 128 characters.
	• The path must include only ASCII characters.
	• The path cannot include non-displayable control characters.
	• The path cannot include any of the following characters:
	`% | < > ' " $ & ; *`

IBM Installation Manager installation directory	The following path is the default path for the IBM Installation Manager installation directory: • `/opt/IBM/InstallationManager`
Port number that is used by the Operations Center web server.	The value for the secure (https) port number must meet the following criteria: • The number must be an integer in the range 1024 - 65535. • The number cannot be in use or allocated to other programs. If you do not specify a port number, the default value is `11090`. **Tip:** If you later do not remember the port number that you specified, refer to the following file, where *`installation_dir`* represents the directory where the Operations Center is installed: • *`installation_dir`*`/ui/Liberty/usr/servers/guiServer/bootstrap.properties` The `bootstrap.properties` file contains the IBM Spectrum Protect server connection information.

Table 25. Information to verify or determine before you install the Operations Center (continued)

<table>
<tr><th>Information</th><th>Details</th></tr>
<tr><td>Password for secure communications</td><td>The Operations Center uses Hypertext Transfer Protocol Secure (HTTPS) to communicate with web browsers.
The Operations Center requires secure communication between the server and the Operations Center. To secure communication, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.
The truststore file of the Operations Center contains the certificate that the Operations Center uses for HTTPS communication with web browsers. During installation of the Operations Center, you create a password for the truststore file. When you set up secure communication between the Operations Center and the hub server, you must use the same password to add the certificate of the hub server to the truststore file.
The password for the truststore file must meet the following criteria:
• The password must contain a minimum of 6 characters and a maximum of 64 characters.
• The password must contain at least the following characters:
– One uppercase letter (A – Z)
– One lowercase letter (a – z)
– One digit (0 – 9)
– Two of the non-alphanumeric characters that are listed in the following series:
~ ! @ # $ % ^ & * _ - + = ` |
( ) { } [ ] : ; < > , . ? /</td></tr>
</table>

Related tasks:

"Configuring for secure communication" on page 133

"Resetting the password for the Operations Center truststore file" on page 136

Chapter 10. Installing the Operations Center

You can install the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

Before you begin

You cannot configure the Operations Center until you install, configure, and start the IBM Spectrum Protect server. Therefore, before you install the Operations Center, install the appropriate server package, according to the server version requirements in "Hub and spoke server requirements" on page 110.

You can install the Operations Center on a computer with the IBM Spectrum Protect server or on a separate computer.

Obtaining the Operations Center installation package

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

About this task

After you obtain the package from an IBM download site, you must extract the installation files.

Procedure

Complete the following steps to extract the Operations Center installation files. In the following steps, replace *version_number* with the version of Operations Center that you are installing.
On Linux systems:

1. Download one of the following package files to the directory of your choice:
   - `version_number.000-IBM-SPOC-LinuxS390.bin`
   - `version_number.000-IBM-SPOC-Linuxx86_64.bin`
2. Ensure that you have executable permission for the package file.

   If necessary, change the file permissions by issuing the following command:

   ```
   chmod a+x package_name.bin
   ```
3. Issue the following command to extract the installation files:

   ```
   ./package_name.bin
   ```

   The self-extracting package file is extracted to the directory.

Installing the Operations Center by using a graphical wizard

You can install or update the Operations Center by using the graphical wizard of IBM Installation Manager.

Procedure

1. From the directory where the Operations Center installation package file is extracted, issue the following command:
   ```
   ./install.sh
   ```
2. Follow the wizard instructions to install the IBM Installation Manager and Operations Center packages.

What to do next

See "Configuring the Operations Center" on page 127.

Installing the Operations Center in console mode

You can install or update the Operations Center by using the command line in console mode.

Procedure

1. From the directory where the installation package file is extracted, run the following program:
   ```
   ./install.sh -c
   ```
2. Follow the console instructions to install the Installation Manager and Operations Center packages.

What to do next

See "Configuring the Operations Center" on page 127.

Installing the Operations Center in silent mode

You can install or upgrade the Operations Center in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the `input` directory where the installation package is extracted:

`install_response_sample.xml`
Use this file to install the Operations Center.

`update_response_sample.xml`
Use this file to upgrade the Operations Center.

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.

   Tip: To generate a response file as part of a console-mode installation, complete the selection of the console-mode installation options. Then, in the Summary panel, enter `G` to generate the response file according to the previously selected options.
2. Create a password for the Operations Center truststore in the response file.

   If you are using the `install_response_sample.xml` file, add the password in the following line of the file, where *mypassword* represents the password:

   ```
   <variable name='ssl.password' value='mypassword' />
   ```

   For more information about this password, see "Installation checklist" on page 118.

   Tip: To upgrade the Operations Center, the truststore password is not required if you are using the `update_response_sample.xml` file.
3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value *response_file* represents the response file path and file name:
   -
     ```
     ./install.sh -s -input response_file -acceptLicense
     ```

What to do next

See "Configuring the Operations Center" on page 127.

Chapter 11. Upgrading the Operations Center

You can upgrade the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

Before you begin

Before you upgrade the Operations Center, review the system requirements and the installation checklist. The new version of the Operations Center might have more or different requirements and considerations than the version you are currently using.

About this task

The instructions for upgrading the Operations Center are the same as the instructions for installing the Operations Center, with the following exceptions:

- You use the **Update** function of IBM Installation Manager rather than the **Install** function.

  Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.
- If you are upgrading the Operations Center in silent mode, you can skip the step of creating a password for the truststore file.

Chapter 12. Getting started with the Operations Center

Before you can use the Operations Center to manage your storage environment, you must configure it.

About this task

After you install the Operations Center, complete the following basic configuration steps:

1. Designate the hub server.
2. Add any spoke servers.
3. Optionally, configure email alerts on the hub and spoke servers.

Figure 1 illustrates an Operations Center configuration.

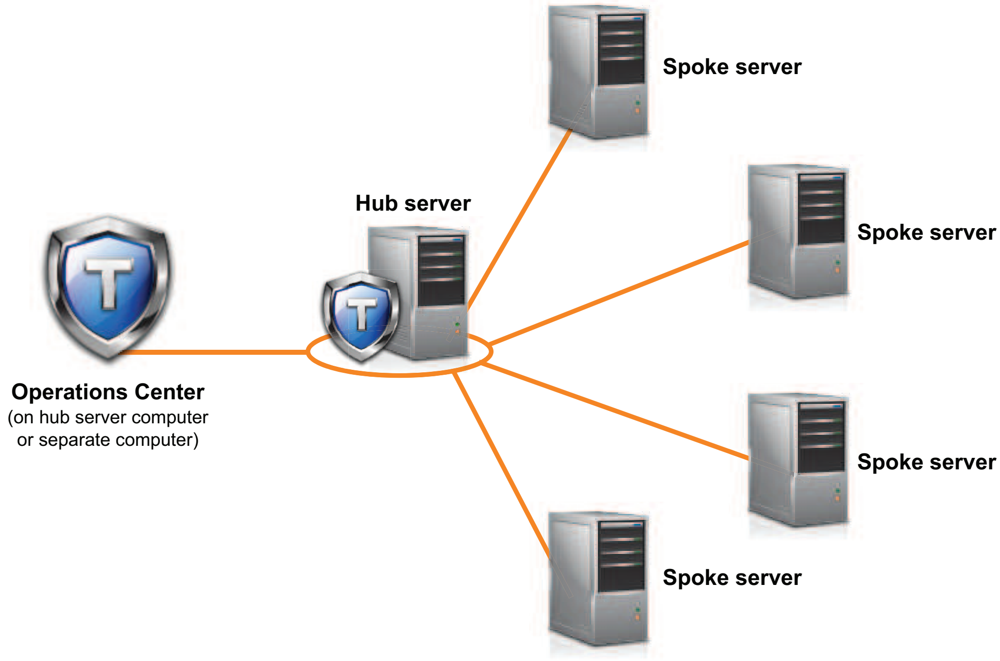

Figure 1. Example of an Operations Center configuration with the hub and spoke servers

Configuring the Operations Center

When you open the Operations Center for the first time, you must configure it to manage your storage environment. You must associate the Operations Center with the IBM Spectrum Protect server that is designated as the hub server. You can then connect additional IBM Spectrum Protect servers as spoke servers.

Designating the hub server

When you connect to the Operations Center for the first time, you must designate which IBM Spectrum Protect server is the hub server.

Before you begin

The Operations Center requires secure communication between the hub server and the Operations Center. To secure communication, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center. For more information, see "Securing communication between the Operations Center and the hub server" on page 133.

Procedure

In a web browser, enter the following address, where *hostname* represents the name of the computer where the Operations Center is installed, and *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer:

```
https://hostname:secure_port/oc
```

Tips:

- The URL is case-sensitive. For example, ensure that you type "oc" in lowercase as indicated.
- For more information about the port number, see the Installation checklist.
- If you are connecting to the Operations Center for the first time, you must provide the following information:
  - Connection information for the server that you want to designate as a hub server
  - Login credentials for an administrator ID that is defined for that server
- If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a hub server.

What to do next

If you have multiple IBM Spectrum Protect servers in your environment, add the other servers as spoke servers to the hub server.

Attention: Do not change the name of a server after it is configured as a hub or spoke server.

Related concepts:

"Hub and spoke server requirements" on page 110

"Administrator IDs that the Operations Center requires" on page 117

Adding a spoke server

After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.

Before you begin

Communication between the spoke server and the hub server must be secured by using the Transport Layer Security (TLS) protocol. To secure communication, add the certificate of the spoke server to the truststore file of the hub server.

Procedure

1. In the Operations Center menu bar, click **Servers**. The Servers page opens.

   In the table on the Servers page, a server might have a status of "Unmonitored." This status means that although an administrator defined this server to the hub server by using the `DEFINE SERVER` command, the server is not yet configured as a spoke server.
2. Complete one of the following steps:
   - Click the server to highlight it, and in the table menu bar, click **Monitor Spoke**.
   - If the server that you want to add is not shown in the table, and secure SSL/TLS communication is not required, click **+ Spoke** in the table menu bar.
3. Provide the necessary information, and complete the steps in the spoke configuration wizard.

   Tip: If the event-record retention period of the server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a spoke server.

Sending email alerts to administrators

An alert is a notification of a relevant problem on the IBM Spectrum Protect server and is triggered by a server message. Alerts can be shown in the Operations Center and can be sent from the server to administrators by email.

Before you begin

Before you configure email notification for administrators about alerts, ensure that the following requirements are met:

- An SMTP server is required to send and receive alerts by email, and the server that sends the alerts by email must have access to the SMTP server.

  Tip: If the Operations Center is installed on a separate computer, that computer does not need access to the SMTP server.
- An administrator must have system privilege to configure email notification.

About this task

An email notification is sent only for the first occurrence of an alert. Also, if an alert is generated before you configure email notification, no email notification is sent for that alert.

You can configure email notification in the following ways:

- Send notification for individual alerts

- Send alert summaries

An alert summary contains information about current alerts. The summary includes the total number of alerts, the total number of active and inactive alerts, the oldest alert, the newest alert, and the most frequently occurring alert.

You can specify a maximum of three administrators to receive alert summaries by email. Alert summaries are sent approximately every hour.

Procedure

To configure email notification for administrators about alerts, complete the following steps on each hub and spoke server from which you want to receive email alerts:

1. To verify that alert monitoring is turned on, issue the following command:
   ```
   QUERY MONITORSETTINGS
   ```
2. If the command output indicates that alert monitoring is turned off, issue the following command. Otherwise, proceed to the next step.
   ```
   SET ALERTMONITOR ON
   ```
3. To enable the sending of email notification, issue the following command:
   ```
   SET ALERTEMAIL ON
   ```
4. To define the SMTP server that is used to send email notification, issue the following command:
   ```
   SET ALERTEMAILSMTPHOST host_name
   ```
5. To specify the port number for the SMTP server, issue the following command:
   ```
   SET ALERTEMAILSMTPPORT port_number
   ```
   The default port number is 25.
6. To specify the email address of the sender of the alerts, issue the following command:
   ```
   SET ALERTEMAILFROMADDR email_address
   ```
7. For each administrator ID that must receive email notification, issue one of the following commands to activate email notification and to specify the email address:
   ```
   REGISTER ADMIN admin_name ALERT=YES EMAILADDRESS=email_address
   UPDATE ADMIN admin_name ALERT=YES EMAILADDRESS=email_address
   ```
8. Choose either, or both, of the following options, and specify the administrator IDs to receive email notification:
   - Send notification for individual alerts

     To specify or update the administrator IDs to receive email notification for an individual alert, issue one of the following commands:
     ```
     DEFINE ALERTTRIGGER message_number ADmin=admin_name1,admin_name2
     UPDATE ALERTTRIGGER message_number ADDadmin=admin_name3 DELadmin=admin_name1
     ```

     Tip: From the Configure Alerts page of the Operations Center, you can select the administrators who will receive email notification.
   - Send alert summaries

     To specify or update the administrator IDs to receive alert summaries by email, issue the following command:
     ```
     SET ALERTSUMMARYTOADMINS admin_name1,admin_name2,admin_name3
     ```

   If you want to receive alert summaries but do not want to receive notification about individual alerts, complete the following steps:

a. Suspend notification about individual alerts, as described in "Suspending email alerts temporarily."

b. Ensure that the respective administrator ID is listed in the following command:

   ```
   SET ALERTSUMMARYTOADMINS admin_name1,admin_name2,admin_name3
   ```

Sending email alerts to multiple administrators

The following example illustrates the commands that cause any alerts for message ANR1075E to be sent in an email to the administrators `myadmin`, `djadmin`, and `csadmin`:

```
SET ALERTMONITOR ON
SET ALERTEMAIL ON
SET ALERTEMAILSMTPHOST mymailserver.domain.com
SET ALERTEMAILSMTPPORT 450
SET ALERTEMAILFROMADDR srvadmin@mydomain.com
UPDATE ADMIN myadmin ALERT=YES EMAILADDRESS=myaddr@anycompany.com
UPDATE ADMIN djadmin ALERT=YES EMAILADDRESS=djaddr@anycompany.com
UPDATE ADMIN csadmin ALERT=YES EMAILADDRESS=csaddr@anycompany.com
DEFINE ALERTTRIGGER anr0175e ADMIN=myadmin,djadmin,csadmin
```

Suspending email alerts temporarily

In certain situations, you might want to suspend email alerts temporarily. For example, you might want to receive alert summaries but suspend notification about individual alerts, or you might want to suspend email alerts when an administrator is on vacation.

Before you begin

Configure email notification for administrators, as described in "Sending email alerts to administrators" on page 129.

Procedure

Suspend email notification for individual alerts or for alert summaries.

- Suspend notification about individual alerts

  Use either of the following methods:

  `UPDATE ADMIN` command

  To turn off email notification for the administrator, issue the following command:

  ```
  UPDATE ADMIN admin_name ALERT=NO
  ```

  To turn on email notification again later, issue the following command:

  ```
  UPDATE ADMIN admin_name ALERT=YES
  ```

  `UPDATE ALERTTRIGGER` command

  To prevent a specific alert from being sent to an administrator, issue the following command:

  ```
  UPDATE ALERTTRIGGER message_number DELADMIN=admin_name
  ```

  To start sending that alert to the administrator again, issue the following command:

  ```
  UPDATE ALERTTRIGGER message_number ADDADMIN=admin_name
  ```

- Suspend notification about alert summaries

To prevent alert summaries from being sent to an administrator, remove the administrator ID from the list in the following command:

```
SET ALERTSUMMARYTOADMINS admin_name1,admin_name2,admin_name3
```

If an administrator ID is listed in the preceding command, the administrator receives alert summaries by email, even if notification about individual alerts is suspended for the respective administrator ID.

Adding customized text to the login screen

You can add customized text, such as your organization's Terms of Use of the software, to the login screen of the Operations Center so that users of the Operations Center see the text before they enter their user name and password.

Procedure

To add customized text to the login screen, complete the following steps:

1. On the computer where the Operations Center is installed, go to the following directory, where `installation_dir` represents the directory in which the Operations Center is installed:

   `installation_dir/ui/Liberty/usr/servers/guiServer`
2. In the directory, create a file that is named `loginText.html` that contains the text that you want to add to the login screen. Any special, non-ASCII text must be UTF-8 encoded.

   Tip: You can format the text by adding HTML tags.
3. Review the added text on the login screen of the Operations Center.

   To open the Operations Center, enter the following address in a web browser, where *hostname* represents the name of the computer where the Operations Center is installed, and *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer:

   `https://hostname:secure_port/oc`

Enabling REST services

Applications that use Representational State Transfer (REST) services can query and manage the storage environment by connecting to the Operations Center.

About this task

Enable this feature to allow REST services to interact with hub and spoke servers by sending calls to the following address:

`https://oc_host_name:port/oc/api`

where *oc_host_name* is the network name or IP address of the Operations Center host system and *port* is the Operations Center port number. The default port number is 11090.

For information about the REST services that are available for the Operations Center, see Technote http://www.ibm.com/support/docview.wss?uid=swg21973011, or issue the following REST call:

`https://oc_host_name:port/oc/api/help`

Procedure

1. On the Operations Center menu bar, hover over the settings icon ⚙ and click **Settings**.
2. On the General page, select the **Enable administrative REST API** check box.
3. Click **Save**.

Configuring for secure communication

The Operations Center uses Hypertext Transfer Protocol Secure (HTTPS) to communicate with web browsers. The Transport Layer Security (TLS) protocol secures communications between the Operations Center and the hub server, and between the hub server and associated spoke servers.

About this task

TLS 1.2 is required for secure communication between the IBM Spectrum Protect server and the Operations Center, and between the hub server and spoke servers.

Securing communication between the Operations Center and the hub server

To secure communications between the Operations Center and the hub server, you must add the Transport Layer Security (TLS) certificate of the hub server to the truststore file of the Operations Center.

Before you begin

The truststore file of the Operations Center is a container for certificates that the Operations Center can access. The truststore file contains the certificate that the Operations Center uses for HTTPS communication with web browsers.

During the installation of the Operations Center, you create a password for the truststore file. To secure communication between the Operations Center and the hub server, you must use the same password to add the certificate of the hub server to the truststore file. If you do not remember this password, you can reset it. See “Resetting the password for the Operations Center truststore file” on page 136.

Procedure

1. Specify the `cert256.arm` certificate as the default certificate in the key database file of the hub server.

   To specify `cert256.arm` as the default certificate, complete the following steps:

   a. Issue the following command from the hub server instance directory:

   ```
   gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed
    -label "TSM Server SelfSigned SHA Key"
   ```

   b. Restart the hub server so that it can receive the changes to the key database file.
2. To verify that the `cert256.arm` certificate is set as the default certificate in the key database file of the hub server, issue the following command:

   ```
   gsk8capicmd_64 -cert -list -db cert.kdb -stashed
   ```
3. Stop the Operations Center web server.
4. Go to the command line of the operating system on which the Operations Center is installed.

5. Add the certificate to the truststore file of the Operations Center by using the **iKeycmd** command or the **iKeyman** command. The **iKeyman** command opens the IBM Key Management graphical user interface, and **iKeycmd** is a command line interface.

   To add the SSL certificate by using the command line interface, issue the **iKeycmd** command to add the `cert256.arm` certificate as the default certificate in the key database file of the hub server:

   ```
   ikeycmd -cert -add
   -db /installation_dir/Liberty/usr/servers/guiServer/gui-truststore.jks
   -file /fvt/comfrey/srv/cert256.arm
   -label 'label description'
   -pw 'password' -type jks -format ascii -trust enable
   ```

   where:

   `installation_dir`
   The directory in which the Operations Center is installed.

   `label description`
   The description that you assign to the label.

   `password`
   The password that you created when you installed the Operations Center. To reset the password, uninstall the Operations Center, delete the `.jks` file, and reinstall the Operations Center.

   To add the certificate by using the IBM Key Management window, complete the following steps:

   a. Go to the following directory, where *`installation_dir`* represents the directory in which the Operations Center is installed:
      - *`installation_dir`*`/ui/jre/bin`
   b. Open the IBM Key Management window by issuing the following command:

      ```
      ikeyman
      ```
   c. Click **Key Database File** > **Open**.
   d. In the Open window, click **Browse**, and go to the following directory, where *`installation_dir`* represents the directory in which the Operations Center is installed:
      - *`installation_dir`*`/ui/Liberty/usr/servers/guiServer`
   e. In the `guiServer` directory, select the `gui-truststore.jks` file.
   f. Click **Open**, and click **OK**.
   g. Enter the password for the truststore file, and click **OK**.
   h. In the **Key database content** area of the IBM Key Management window, click the arrow, and select **Signer Certificates** from the list.
   i. Click **Add**.
   j. In the Open window, click **Browse**, and go to the hub server instance directory, as shown in the following example:
      - `/opt/tivoli/tsm/server/bin`

      The directory contains the `cert256.arm` certificate.

      If you cannot access the hub server instance directory from the Open window, complete the following steps:

      1) Use FTP or another file-transfer method to copy the `cert256.arm` files from the hub server to the following directory on the computer where the Operations Center is installed:

    - `installation_dir/ui/Liberty/usr/servers/guiServer`
  2) In the Open window, go to the `guiServer` directory.

k. Select the `cert256.arm` certificate as the certificate.

   Tip: The certificate that you choose must be set as the default certificate in the key database file of the hub server. For more information, see step 1 on page 133 and 2 on page 133.

l. Click **Open**, and click **OK**.

m. Enter a label for the certificate. For example, enter the name of the hub server.

n. Click **OK**. The SSL certificate of the hub server is added to the truststore file, and the label is displayed in the **Key database content** area of the IBM Key Management window.

o. Close the IBM Key Management window.

6. Start the Operations Center web server.
7. When you connect to the Operations Center for the first time, you are prompted to identify the IP address or network name of the hub server, and the port number for communicating with the hub server. If the ADMINONCLIENTPORT server option is enabled for the IBM Spectrum Protect server, enter the port number that is specified by the TCPADMINPORT server option. If the ADMINONCLIENTPORT server option is not enabled, enter the port number that is specified by the TCPPORT server option.

   If the Operations Center was previously configured, you can review the contents of the `serverConnection.properties` file to verify the connection information. The `serverConnection.properties` file is in the following directory on the computer where the Operations Center is installed:

   - `installation_dir/ui/Liberty/usr/servers/guiServer`

What to do next

To set up SSL communication between the hub server and a spoke server, see "Securing communication between the hub server and a spoke server."

Securing communication between the hub server and a spoke server

To secure communications between the hub server and a spoke server by using the Transport Layer Security (TLS) protocol, you must define the certificate of the spoke server to the hub server. You must also configure the Operations Center to monitor the spoke server.

Procedure

1. On the spoke server, change to the directory of the spoke server instance.
2. Specify the required `cert256.arm` certificate as the default certificate in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed
 -label "TSM Server SelfSigned SHA Key"
```

3. Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

4. Securely transfer the `cert256.arm` file of the spoke server to the hub server.
5. On the hub server, change to the directory of the hub server instance.

6. Define the spoke server certificate to the hub server. Issue the following command from the hub server instance directory, where *spoke_servername* is the name of the spoke server, and *spoke_cert256.arm* is the file name of the spoke server certificate:

   ```
   gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii
    -label spoke_servername -file spoke_cert256.arm
   ```

   The spoke server does not require the hub server certificate for hub-to-spoke communication. However, other server configurations that require cross-defined servers do require the spoke server to have the hub server certificate.
7. Restart the hub server and the spoke server.
8. For the hub server, issue the **DEFINE SERVER** command, according to the following example:

   ```
   DEFINE SERVER spoke_servername HLA=spoke_address
    LLA=spoke_SSLTCPADMINPort SERVERPA=spoke_serverpassword
   ```

   Tip: By default, server communication is encrypted except when the server is sending or receiving object data. Object data is sent and received by using TCP/IP. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP session and the session is secure. To encrypt all communication with the specified server, even when the server is sending and receiving object data, specify the SSL=YES parameter on the **DEFINE SERVER** command.
9. On the Operations Center menu bar, click **Servers**.

   In the table on the Servers page, the spoke server that you defined in step 8 typically has a status of "Unmonitored." Depending on the setting for the status refresh interval, you might not see the spoke server immediately.
10. Click the spoke server to highlight the item, and in the table menu bar, click **Monitor Spoke**.

Resetting the password for the Operations Center truststore file

To set up secure communication between the Operations Center and the hub server, you must know the password for the truststore file of the Operations Center. You create this password during the installation of the Operations Center. If you do not know the password, you can reset it.

About this task

To reset the password, you must create a new password, delete the truststore file of the Operations Center, and restart the Operations Center web server.

Procedure

1. Stop the Operations Center web server.
2. Go to the following directory, where `installation_dir` represents the directory in which the Operations Center is installed:

   ```
   installation_dir/ui/Liberty/usr/servers/guiServer
   ```
3. Open the `bootstrap.properties` file, which contains the password for the truststore file. If the password is unencrypted, you can use it to open the truststore file without having to reset it.

   The following examples indicate the difference between an encrypted and an unencrypted password:

Encrypted password example

Encrypted passwords begin with the text string `{xor}`.

The following example shows the encrypted password as the value of the **`tsm.truststore.pswd`** parameter:

```
tsm.truststore.pswd={xor}MiYPPiwsKDAtOw==
```

Unencrypted password example

The following example shows the unencrypted password as the value of the **`tsm.truststore.pswd`** parameter:

```
tsm.truststore.pswd=J8b%^B
```

4. Reset the password by replacing the password in the `bootstrap.properties` file with a new password. You can replace the password with an encrypted or unencrypted password. Remember the unencrypted password for future use.

   To create an encrypted password, complete the following steps:

   a. Create an unencrypted password.

      The password for the truststore file must meet the following criteria:

      - The password must contain a minimum of 6 characters and a maximum of 64 characters.
      - The password must contain at least the following characters:
        - One uppercase letter (A – Z)
        - One lowercase letter (a – z)
        - One digit (0 – 9)
        - Two of the non-alphanumeric characters that are listed in the following series:

          `~ ! @ # $ % ^ & * _ - + = ` |`

          `( ) { } [ ] : ; < > , . ? /`

   b. From the command line of the operating system, go to the following directory:

      ```
      installation_dir/ui/Liberty/bin
      ```

   c. To encrypt the password, issue the following command, where *`myPassword`* represents the unencrypted password:

      ```
      securityUtility encode myPassword
      ```

5. Close the `bootstrap.properties` file.
6. Go to the following directory:

   ```
   installation_dir/ui/Liberty/usr/servers/guiServer
   ```

7. Delete the `gui-truststore.jks` file, which is the truststore file of the Operations Center.
8. Start the Operations Center web server.

Results

A new truststore file is automatically created for the Operations Center, and the TLS certificate of the Operations Center is automatically included in the truststore file.

Starting and stopping the web server

The web server of the Operations Center runs as a service and starts automatically. You might need to stop and start the web server, for example, to make configuration changes.

Procedure

Stop and start the web server.

- Issue the following commands:
  - To stop the server:
    ```
    service opscenter.rc stop
    ```
  - To start the server:
    ```
    service opscenter.rc start
    ```
  - To restart the server:
    ```
    service opscenter.rc restart
    ```

  To determine whether the server is running, issue the following command:

  ```
  service opscenter.rc status
  ```

Opening the Operations Center

The Overview page is the default initial view in the Operations Center. However, in your web browser, you can bookmark the page that you want to open when you log in to the Operations Center.

Procedure

1. In a web browser, enter the following address, where *hostname* represents the name of the computer where the Operations Center is installed, and *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer:

   ```
   https://hostname:secure_port/oc
   ```

   Tips:

   - The URL is case-sensitive. For example, ensure that you type “oc” in lowercase as indicated.
   - The default port number for HTTPS communication is 11090, but a different port number can be specified during Operations Center installation.
2. Log in, using an administrator ID that is registered on the hub server.

   In the Overview page, you can view summary information for clients, services, servers, storage pools, and storage devices. You can view more details by clicking items or by using the Operations Center menu bar.

   Monitoring from a mobile device: To remotely monitor the storage environment, you can view the Overview page of the Operations Center in the web browser of a mobile device. The Operations Center supports the Apple Safari web browser on the iPad. Other mobile devices can also be used.

Collecting diagnostic information with IBM Spectrum Protect client management services

The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

About this task

After you install the client management service, you can view the Diagnosis page in the Operations Center to obtain troubleshooting information for backup-archive clients.

Diagnostic information can be collected only from Linux and Windows clients, but administrators can view the diagnostic information in the Operations Center on AIX, Linux, or Windows operating systems.

You can also install the client management service on data mover nodes for IBM Spectrum Protect for Virtual Environments: Data Protection for VMware to collect diagnostic information about the data movers.

Tip: In the documentation for the client management service, *client system* is the system where the backup-archive client is installed.

Installing the client management service by using a graphical wizard

To collect diagnostic information about backup-archive clients such as client log files, you must install the client management service on the client systems that you manage.

Before you begin

Review "Requirements and limitations for IBM Spectrum Protect client management services" on page 115.

About this task

You must install the client management service on the same computer as the backup-archive client.

Procedure

1. Download the installation package for the client management service from an IBM download site such as IBM Passport Advantage or IBM Fix Central. Look for a file name that is similar to `<`*`version`*`>-IBM-SPCMS-<`*`operating system`*`>.bin`.

   The following table shows the names of the installation packages.

Client operating system	**Installation package name**
Linux x86 64-bit	`8.1.x.000-IBM-SPCMS-Linuxx64.bin`
Windows 32-bit	`8.1.x.000-IBM-SPCMS-Windows32.exe`
Windows 64-bit	`8.1.x.000-IBM-SPCMS-Windows64.exe`

2. Create a directory on the client system that you want to manage, and copy the installation package there.

3. Extract the contents of the installation package file.
   - On Linux client systems, complete the following steps:
     a. Change the file to an executable file by issuing the following command:
        ```
        chmod +x 8.1.x.000-IBM-SPCMS-Linuxx64.bin
        ```
     b. Issue the following command:
        ```
        ./8.1.x.000-IBM-SPCMS-Linuxx64.bin
        ```
   - On Windows client systems, double-click the installation package name in Windows Explorer.

   Tip: If you previously installed and uninstalled the package, select **All** when prompted to replace the existing installation files.
4. Run the installation batch file from the directory where you extracted the installation files and associated files. This is the directory that you created in step 2 on page 139.
   - On Linux client systems, issue the following command:
     ```
     ./install.sh
     ```
   - On Windows client systems, double-click **install.bat**.
5. To install the client management service, follow the instructions in the IBM Installation Manager wizard.

   If IBM Installation Manager is not already installed on the client system, you must select both **IBM Installation Manager** and **IBM Spectrum Protect Client Management Services**.

   Tip: You can accept the default locations for the shared resources directory and the installation directory for IBM Installation Manager.

What to do next

Follow the instructions in “Verifying that the client management service is installed correctly” on page 141.

Installing the client management service in silent mode

You can install the client management service in silent mode. When you use silent mode, you provide the installation values in a response file and then run an installation command.

Before you begin

Review “Requirements and limitations for IBM Spectrum Protect client management services” on page 115.

Extract the installation package by following the instructions in “Installing the client management service by using a graphical wizard” on page 139.

About this task

You must install the client management service on the same computer as the backup-archive client.

The `input` directory, which is in the directory where the installation package is extracted, contains the following sample response file:

```
install_response_sample.xml
```

You can use the sample file with the default values, or you can customize it.

Tip: If you want to customize the sample file, create a copy of the sample file, rename it, and edit the copy.

Procedure

1. Create a response file based on the sample file, or use the sample file, `install_response_sample.xml`.

   In either case, ensure that the response file specifies the port number for the client management service. The default port is 9028. For example:

   ```
   <variable name='port' value='9028'/>
   ```
2. Run the command to install the client management service and accept the license. From the directory where the installation package file is extracted, issue the following command, where *response_file* represents the response file path, including the file name:

   On a Linux client system:

   ```
   ./install.sh -s -input response_file -acceptLicense
   ```

   For example:

   ```
   ./install.sh -s -input /cms_install/input/install_response.xml -acceptLicense
   ```

   On a Windows client system:

   ```
   install.bat -s -input response_file -acceptLicense
   ```

   For example:

   ```
   install.bat -s -input c:\cms_install\input\install_response.xml -acceptLicense
   ```

What to do next

Follow the instructions in “Verifying that the client management service is installed correctly.”

Verifying that the client management service is installed correctly

Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.

Procedure

On the client system, at the command line, run the following commands to view the configuration of the client management service:

- On Linux client systems, issue the following command:

  ```
  client_install_dir/cms/bin/CmsConfig.sh list
  ```

  where *client_install_dir* is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

  ```
  /opt/tivoli/tsm/cms/bin/CmsConfig.sh list
  ```

  The output is similar to the following text:

  ```
  Listing CMS configuration

  server1.example.com:1500 NO_SSL HOSTNAME
  Capabilities: [LOG_QUERY]
  ```

```
    Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys

    Log File: /opt/tivoli/tsm/client/ba/bin/dsmerror.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252

    Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

- On Windows client systems, issue the following command:

  ```
  client_install_dir\cms\bin\CmsConfig.bat list
  ```

  where *client_install_dir* is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

  ```
  C:\"Program Files"\Tivoli\TSM\cms\bin\CmsConfig.bat list
  ```

  The output is similar to the following text:

  ```
  Listing CMS configuration

  server1.example.com:1500 NO_SSL HOSTNAME
  Capabilities: [LOG_QUERY]
      Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt

      Log File: C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
          en_US MM/dd/yyyy HH:mm:ss Windows-1252

      Log File: C:\Program Files\Tivoli\TSM\baclient\dsmsched.log
          en_US MM/dd/yyyy HH:mm:ss Windows-1252
  ```

If the client management service is correctly installed and configured, the output displays the location of the error log file.
The output text is extracted from the following configuration file:

- On Linux client systems:

  ```
  client_install_dir/cms/Liberty/usr/servers/cmsServer/client-configuration.xml
  ```

- On Windows client systems:

  ```
  client_install_dir\cms\Liberty\usr\servers\cmsServer\client-configuration.xml
  ```

If the output does not contain any entries, you must configure the `client-configuration.xml` file. For instructions about how to configure this file, see "Configuring the client management service for custom client installations" on page 144. You can use the **`CmsConfig verify`** command to verify that a node definition is correctly created in the `client-configuration.xml` file.

Configuring the Operations Center to use the client management service

If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Before you begin

Ensure that the client management service is installed and started on the client system.

Verify whether the default configuration is used. The default configuration is not used if either of the following conditions is met:

- The client management service does not use the default port number, 9028.

- The backup-archive client is not accessed by the same IP address as the client system where the backup-archive client is installed. For example, a different IP address might be used in the following situations:
  - The computer system has two network cards. The backup-archive client is configured to communicate on one network, while the client management service communicates on the other network.
  - The client system is configured with the Dynamic Host Configuration Protocol (DHCP). As a result, the client system is dynamically assigned an IP address, which is saved on the IBM Spectrum Protect server during the previous backup-archive client operation. When the client system is restarted, the client system might be assigned a different IP address. To ensure that the Operations Center can always find the client system, you specify a fully qualified domain name.

Procedure

To configure the Operations Center to use the client management service, complete the following steps:

1. On the Clients page of the Operations Center, select the client.
2. Click **Details**.
3. Click the **Properties** tab.
4. In the **Remote diagnostics URL** field in the **General** section, specify the URL for the client management service on the client system.

   The address must start with `https`. The following table shows examples of the remote diagnostics URL.

Type of URL	Example
With DNS host name and default port, 9028	`https://server.example.com`
With DNS host name and non-default port	`https://server.example.com:1599`
With IP address and non-default port	`https://192.0.2.0:1599`

5. Click **Save**.

What to do next

You can access client diagnostic information such as client log files from the **Diagnosis** tab in the Operations Center.

Starting and stopping the client management service

The client management service is automatically started after it is installed on the client system. You might need to stop and start the service in certain situations.

Procedure

- To stop, start, or restart the client management service on Linux client systems, issue the following commands:
  - To stop the service:

    ```
    service cms.rc stop
    ```

  - To start the service:

    ```
    service cms.rc start
    ```

  - To restart the service:

    ```
    service cms.rc restart
    ```

- On Windows client systems, open the Services window, and stop, start, or restart the IBM Spectrum Protect Client Management Services service.

Uninstalling the client management service

If you no longer have to collect client diagnostic information, you can uninstall the client management service from the client system.

About this task

You must use IBM Installation Manager to uninstall the client management service. If you no longer plan to use IBM Installation Manager, you can also uninstall it.

Procedure

1. Uninstall the client management service from the client system:
   a. Open IBM Installation Manager:
      - On the Linux client system, in the directory where IBM Installation Manager is installed, go to the `eclipse` subdirectory (for example, `/opt/IBM/InstallationManager/eclipse`), and issue the following command:
        ```
        ./IBMIM
        ```
      - On the Windows client system, open IBM Installation Manager from the **Start** menu.
   b. Click **Uninstall**.
   c. Select **IBM Spectrum Protect Client Management Services**, and click **Next**.
   d. Click **Uninstall**, and then click **Finish**.
   e. Close the IBM Installation Manager window.
2. If you no longer require IBM Installation Manager, uninstall it from the client system:
   a. Open the IBM Installation Manager uninstall wizard:
      - On the Linux client system, change to the IBM Installation Manager uninstallation directory (for example, `/var/ibm/InstallationManager/uninstall`), and issue the following command:
        ```
        ./uninstall
        ```
      - On the Windows client system, click **Start** > **Control Panel**. Then, click **Uninstall a program** > **IBM Installation Manager** > **Uninstall**.
   b. In the IBM Installation Manager window, select **IBM Installation Manager** if it is not already selected, and click **Next**.
   c. Click **Uninstall**, and click **Finish**.

Configuring the client management service for custom client installations

The client management service uses information in the client configuration file (`client-configuration.xml`) to discover diagnostic information. If the client management service is unable to discover the location of log files, you must run the **`CmsConfig`** utility to add the location of the log files to the `client-configuration.xml` file.

`CmsConfig` **utility**

If you are not using the default client configuration, you can run the **`CmsConfig`** utility on the client system to discover and add the location of the client log files to the `client-configuration.xml` file. After you complete the configuration, the client management service can access the client log files and make them available for basic diagnostic functions in the Operations Center.

You can also use the **`CmsConfig`** utility to show the configuration of the client management service and to remove a node name from the `client-configuration.xml` file.

The `client-configuration.xml` file is in the following directory:

- On Linux client systems:
  `client_install_dir``/cms/Liberty/usr/servers/cmsServer`
- On Windows client systems:
  `client_install_dir``\cms\Liberty\usr\servers\cmsServer`

where *client_install_dir* is the directory where the backup-archive client is installed.

The **`CmsConfig`** utility is available in the following locations.

Client operating system	**Utility location and name**
Linux	*`client_install_dir`*`/cms/bin/CmsConfig.sh`
Windows	*`client_install_dir`*`\cms\bin\CmsConfig.bat`

To use the **`CmsConfig`** utility, issue any command that is included in the utility. Ensure that you enter each command on a single line.

`CmsConfig discover` command:

You can use the **`CmsConfig discover`** command to automatically discover options files and log files, and add them to the client configuration file, `client-configuration.xml`. In this way, you can help to ensure that the client management service can access the client log files and make them available for diagnosis in the Operations Center.

Typically, the client management service installer runs the **`CmsConfig discover`** command automatically. However, you must run this command manually if you changed the backup-archive client, such as added a client, or changed the server configuration or location of log files.

For the client management service to create a log definition in the `client-configuration.xml` file, the IBM Spectrum Protect server address, server port, and client node name must be obtained. If the node name is not defined in the client options file (typically, `dsm.sys` on Linux client systems and `dsm.opt` on Windows client systems), the host name of the client system is used.

To update the client configuration file, the client management service must access one or more log files, such as `dsmerror.log` and `dsmsched.log`. For best results, run the **`CmsConfig discover`** command in the same directory and by using the same environment variables as you would for the backup-archive client command, **`dsmc`**. In this way, you can improve the chances of finding the correct log files.

If the client options file is in a custom location or it does not have a typical options file name, you can also specify the path for the client options file to narrow the scope of the discovery.

Syntax

```
►►──CmsConfig discover──┬────────────┬──────────────────────────────►◄
                        └─configPath─┘
```

Parameters

`configPath`
: The path of the client options file (typically `dsm.opt`). Specify the configuration path when the client options file is not in a default location or it does not have the default name. The client management service loads the client options file and discovers the client nodes and logs from there. This parameter is optional.

On a Linux client system, the client management service always loads the client user-options file (`dsm.opt`) first, and then looks for the client system-options file (typically `dsm.sys`). The value of the *configPath* parameter, however, is always the client user-options file.

Examples for a Linux client system

- Discover the client log files and automatically add the log definitions to the `client-configuration.xml` file.

  Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

  Command:

  ```
  ./CmsConfig.sh discover
  ```

  Output:

  ```
  Discovering client configuration and logs.

  server.example.com:1500 SUSAN
      /opt/tivoli/tsm/client/ba/bin/dsmerror.log

  Finished discovering client configuration and logs.
  ```

- Discover the configuration files and log files that are specified in the `/opt/tivoli/tsm/client/ba/bin/daily.opt` file and automatically add the log definitions to the `client-configuration.xml` file.

  Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

  Command:

  ```
  ./CmsConfig.sh discover /opt/tivoli/tsm/client/ba/bin/daily.opt
  ```

  Output:

  ```
  Discovering client configuration and logs

  server.example.com:1500 NO_SSL SUSAN
  Capabilities: [LOG_QUERY]
      Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys

      Log File:  /opt/tivoli/tsm/client/ba/bin/dsmerror.log
          en_US MM/dd/yyyy HH:mm:ss Windows-1252

      Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
          en_US MM/dd/yyyy HH:mm:ss Windows-1252

  Finished discovering client configuration and logs.
  ```

Examples for a Windows client system

- Discover the client log files and automatically add the log definitions to the `client-configuration.xml` file.

  Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

  Command:

  ```
  cmsconfig discover
  ```

  Output:

  ```
  Discovering client configuration and logs.

  server.example.com:1500 SUSAN
      C:\Program Files\Tivoli\TSM\baclient\dsmerror.log

  Finished discovering client configuration and logs.
  ```
- Discover the configuration files and log files that are specified in the `c:\program files\tivoli\tsm\baclient\daily.opt` file and automatically add the log definitions to the `client-configuration.xml` file.

  Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

  Command:

  ```
  cmsconfig discover "c:\program files\tivoli\tsm\baclient\
  daily.opt"
  ```

  Output:

  ```
  Discovering client configuration and logs

  server.example.com:1500 NO_SSL SUSAN
  Capabilities: [LOG_QUERY]
      Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt

      Log File: C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
          en_US MM/dd/yyyy HH:mm:ss Windows-1252

      Log File: C:\Program Files\Tivoli\TSM\baclient\dsmsched.log
          en_US MM/dd/yyyy HH:mm:ss Windows-1252

  Finished discovering client configuration and logs.
  ```

`CmsConfig addnode` command:

Use the **`CmsConfig addnode`** command to manually add a client node definition to the `client-configuration.xml` configuration file. The node definition contains information that is required by the client management service to communicate with the IBM Spectrum Protect server.

Use this command only if the client options file or client log files are stored in a non-default location on the client system.

Syntax

```
►►—CmsConfig addnode—nodeName—serverIP—serverPort—serverProtocol—optPath———►◄
```

Parameters

nodeName

The client node name that is associated with the log files. For most client systems, only one node name is registered to the IBM Spectrum Protect server. However, on systems with multiple users, such as Linux client systems, there can be more than one client node name. This parameter is required.

serverIP

The TCP/IP address of the IBM Spectrum Protect server that the client management service authenticates to. This parameter is required.

You can specify a 1 - 64 character TCP/IP address for the server. The server address can be a TCP/IP domain name or a numeric IP address. The numeric IP address can be either a TCP/IP v4 or TCP/IP v6 address. You can use IPv6 addresses only if the **`commmethod V6Tcpip`** option is specified for the client system.

Examples:

- `server.example.com`
- `192.0.2.0`
- `2001:0DB8:0:0:0:0:0:0`

serverPort

The TCP/IP port number that is used to communicate with the IBM Spectrum Protect server. You can specify a value in the range 1 - 32767. This parameter is required.

Example: `1500`

serverProtocol

The protocol that is used for communication between the client management service and the IBM Spectrum Protect server. This parameter is required.

You can specify one of the following values.

Value	Meaning
NO_SSL	The SSL security protocol is not used.
SSL	The SSL security protocol is used.
FIPS	The TLS 1.2 protocol is used in Federal Information Processing Standard (FIPS) mode. **Tip:** Alternatively, you can enter `TLS_1.2` to specify that the TLS 1.2 protocol is used in FIPS mode.

optPath

The fully qualified path of the client options file. This parameter is required.

Example (Linux client): `/opt/backup_tools/tivoli/tsm/baclient/dsm.sys`

Example (Windows client): `C:\backup tools\Tivoli\TSM\baclient\dsm.opt`

Example for a Linux client system

Add the node definition for client node `SUSAN` to the `client-configuration.xml` file. The IBM Spectrum Protect server that the node communicates with is `server.example.com` on server port `1500`. The SSL security protocol is not used. The path for the client system options file is `/opt/tivoli/tsm/client/ba/bin/custom_opt.sys`.

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh addnode SUSAN server.example.com 1500 NO_SSL
/opt/tivoli/tsm/client/ba/bin/custom_opt.sys
```

Output:

```
Adding node.

Finished adding client configuration.
```

Example for a Windows client system

Add the node definition for client node `SUSAN` to the `client-configuration.xml` file. The IBM Spectrum Protect server that the node communicates with is `server.example.com` on server port `1500`. The SSL security protocol is not used. The path for the client options file is `c:\program files\tivoli\tsm\baclient\custom.opt`.

Issue the following command. from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

Command:

```
cmsconfig addnode SUSAN server.example.com 1500 NO_SSL "c:\program
files\tivoli\tsm\baclient\custom.opt"
```

Output:

```
Adding node.

Finished adding client configuration.
```

`CmsConfig setopt` command:

Use the **`CmsConfig setopt`** command to set the path of the client options file (typically `dsm.opt`) to an existing node definition without first reading the contents of the client options file.

This command can be helpful if the client options file does not have a typical name or is in a non-default location.

Requirement: If the node definition does not exist, you must first issue the **`CmsConfig addnode`** command to create the node definition.

Unlike the **`CmsConfig discover`** command, the **`CmsConfig setopt`** command does not create associated log definitions in the `client-configuration.xml` file. You must use the **`CmsComfog addlog`** command to create the log definitions.

Syntax

```
►►──CmsConfig setopt──nodeName──optPath──────────────────────────────────►◄
```

Parameters

nodeName
The client node name that is associated with the log files. For most client systems, only one node name is registered to the IBM Spectrum Protect server. However, on systems with multiple users, such as Linux client systems, there can be more than one client node name. This parameter is required.

`optPath`

The fully qualified path of the client options file. This parameter is required.

Example (Linux client): `/opt/backup_tools/tivoli/tsm/baclient/dsm.opt`

Example (Windows client): `C:\backup tools\Tivoli\TSM\baclient\dsm.opt`

Example for a Linux client system

Set the client options file path for the node `SUSAN`. The path for the client options file is `/opt/tivoli/tsm/client/ba/bin/dsm.opt`.

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh setopt SUSAN /opt/tivoli/tsm/client/ba/bin/dsm.opt
```

Output:

```
Adding node configuration file.

Finished adding client configuration file.
```

Example for a Windows client system

Set the client options file path for the node `SUSAN`. The path for the client options file is `c:\program files\tivoli\tsm\baclient\dsm.opt`.

Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

Command:

```
cmsconfig setopt SUSAN "c:\program files\tivoli\tsm\baclient\
dsm.opt"
```

Output:

```
Adding node configuration file.

Finished adding client configuration file.
```

`CmsConfig setsys` command:

On a Linux client system, use the **`CmsConfig setsys`** command to set the path of the client system-options file (typically `dsm.sys`) to an existing node definition without first reading the contents of the client system-options file.

This command can be helpful if the client system-options file does not have a typical name or is in a non-default location.

Requirement: If the node definition does not exist, you must first issue the **`CmsConfig addnode`** command to create the node definition.

Unlike the **`CmsConfig discover`** command, the **`CmsConfig setsys`** command does not create associated log definitions in the `client-configuration.xml` file. You must use the **`CmsComfog addlog`** command to create the log definitions.

Syntax

```
►►──CmsConfig setsys──nodeName──sysPath──────────────────────────────►◄
```

Parameters

nodeName
: The client node name that is associated with the log files. For most client systems, only one node name is registered to the IBM Spectrum Protect server. However, on systems with multiple users, such as Linux client systems, there can be more than one client node name. This parameter is required.

sysPath
: The fully qualified path of the client system-options file. This parameter is required.

Example: `/opt/backup_tools/tivoli/tsm/baclient/dsm.sys`

Example

Set the client system-options file path for the node `SUSAN`. The path for the client system-options file is `/opt/tivoli/tsm/client/ba/bin/dsm.sys`.

Issue the following command, from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh setopt SUSAN /opt/tivoli/tsm/client/ba/bin/dsm.sys
```

Output:

```
Adding node configuration file.

Finished adding client configuration file.
```

`CmsConfig addlog` command:

Use the **`CmsConfig addlog`** command to manually add the location of client log files to an existing node definition in the `client-configuration.xml` configuration file. Use this command only if the client log files are stored in a non-default location on the client system.

Requirement: If the node definition does not exist, you must first issue the **`CmsConfig addnode`** command to create the node definition.

Syntax

```
►►──CmsConfig addlog──nodeName──logPath─────────────────────────────────►

►──┬───────────────────────────────────────────────┬───────────────────►◄
   └─language──dateFormat──timeFormat──encoding────┘
```

Parameters

nodeName
: The client node name that is associated with the log files. For most client systems, only one node name is registered to the IBM Spectrum Protect server. However, on systems with multiple users, such as Linux client systems, there can be more than one client node name. This parameter is required.

logPath
: The fully qualified path of the log files. This parameter is required.

Example (Linux client): `/opt/backup_tools/tivoli/tsm/baclient/dsmerror.log`

Example (Windows client): `C:\backup tools\Tivoli\TSM\baclient\dsmerror.log`

language
The language locale of the log file. This parameter is optional. However, if you specify this parameter, you must also specify the **dateFormat**, **timeFormat**, and **encoding** parameters. You must specify the locale for the following languages.

Language	Locale
Brazilian Portuguese	pt_BR
Chinese, Simplified	zh_CN
Chinese, Traditional	zh_TW
Czech	cs_CZ
English	en_US
French	fr_FR
German	de_DE
Hungarian	hu_HU
Italian	it_IT
Japanese	ja_JP
Korean	ko_KR
Polish	pl_PL
Russian	ru_RU
Spanish	es_ES

dateFormat
The date format of the time stamp entries in the client log file. This parameter is optional. However, if you specify this parameter, you must also specify the **language**, **timeFormat**, and **encoding** parameters.

The following table shows the date formats for the languages.

Tip: Instead of using one of the date formats that are listed in the table, you can specify a date format by using the backup-archive client **dateformat** option.

Language	Date format
Chinese, Simplified	yyyy-MM-dd
Chinese, Traditional	yyyy/MM/dd
Czech	dd.MM.yyyy
English	MM/dd/yyyy
French	dd/MM/yyyy
German	dd.MM.yyyy
Hungarian	yyyy.MM.dd
Italian	dd/MM/yyyy
Japanese	yyyy-MM-dd
Korean	yyyy/MM/dd
Polish	yyyy-MM-dd
Portuguese, Brazilian	dd/MM/yyyy

Language	Date format
Russian	`dd.MM.yyyy`
Spanish	`dd.MM.yyyy`

`timeFormat`

The time format of the time stamp entries in the client log file. This parameter is optional. However, if you specify this parameter, you must also specify the **`language`**, **`dateFormat`**, and **`encoding`** parameters.

The following table shows examples of default time formats that you can specify and client operating systems.

Tip: Instead of using one of the time formats that are listed in the table, you can specify a time format by using the backup-archive client **`timeformat`** option.

Language	Time format for Linux client systems	Time format for Windows client systems
Chinese, Simplified	`HH:mm:ss`	`HH:mm:ss`
Chinese, Traditional	`HH:mm:ss`	`ahh:mm:ss`
Czech	`HH:mm:ss`	`HH:mm:ss`
English	`HH:mm:ss`	`HH:mm:ss`
French	`HH:mm:ss`	`HH:mm:ss`
German	`HH:mm:ss`	`HH:mm:ss`
Hungarian	`HH.mm.ss`	`HH:mm:ss`
Italian	`HH:mm:ss`	`HH:mm:ss`
Japanese	`HH:mm:ss`	`HH:mm:ss`
Korean	`HH:mm:ss`	`HH:mm:ss`
Polish	`HH:mm:ss`	`HH:mm:ss`
Portuguese, Brazilian	`HH:mm:ss`	`HH:mm:ss`
Russian	`HH:mm:ss`	`HH:mm:ss`
Spanish	`HH:mm:ss`	`HH:mm:ss`

`encoding`

The character encoding of the entries in the client log files. This parameter is optional. However, if you specify this parameter, you must also specify the **`language`**, **`dateFormat`**, and **`timeFormat`** parameters.

For Linux client systems, the typical character encoding is UTF-8. For Windows client systems, the default encoding values are shown in the following table. If your client system is customized differently, use the **`encoding`** parameter to specify a value other than the default.

Language	Encoding
Chinese, Simplified	`CP936`
Chinese, Traditional	`CP950`
Czech	`Windows-1250`
English	`Windows-1252`
French	`Windows-1252`

Language	Encoding
German	`Windows-1252`
Hungarian	`Windows-1250`
Italian	`Windows-1252`
Japanese	`CP932`
Korean	`CP949`
Polish	`Windows-1250`
Portuguese, Brazilian	`Windows-1252`
Russian	`Windows-1251`
Spanish	`Windows-1252`

Example for a Linux client system

Add the client log file location to the existing definition for client node SUSAN in the `client-configuration.xml` file. The path for the client log file is `/usr/work/logs/dsmerror.log`. Add the language specification, time format, and date format for the French locale.

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh addlog SUSAN /usr/work/logs/dsmerror.log fr_FR
yyyy/MM/dd HH:MM:ss UTF-8
```

Output:

```
Adding log.

Finished adding log.
```

Example for a Windows client system

Add the client log file location to the existing definition for client node SUSAN in the `client-configuration.xml`. The path for the client log file is `c:\work\logs\dsmerror.log`. Add the language specification, time format, and date format for the French locale.

Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

Command:

```
cmsconfig addlog SUSAN c:\work\logs\dsmerror.log fr_FR yyyy/MM/dd
HH:MM:ss UTF-8
```

Output:

```
Adding log.

Finished adding log.
```

CmsConfig remove command:

Use the **CmsConfig remove** command to remove a client node definition from the client configuration file, `client-configuration.xml`. All log file entries that are associated with the client node name are also removed.

Syntax

```
►►──CmsConfig remove──nodeName──────────────────────────────────────────►◄
```

Parameters

nodeName
The client node name that is associated with the log files. For most client systems, only one node name is registered to the IBM Spectrum Protect server. However, on systems with multiple users, such as Linux client systems, there can be more than one client node name. This parameter is required.

Example for a Linux client system

Remove the node definition for `SUSAN` from the `client-configuration.xml` file.

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh remove SUSAN
```

Output:

```
Removing node.

Finished removing node.
```

Example for a Windows client system

Remove the node definition for `SUSAN` from the `client-configuration.xml` file.

Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

Command:

```
cmsconfig remove SUSAN
```

Output:

```
Removing node.

Finished removing node.
```

CmsConfig verify command:

Use the **CmsConfig verify** command to verify that a node definition is correctly created in the `client-configuration.xml` file. If there are errors with the node definition or the node is not correctly defined, you must correct the node definition by using the appropriate **CmsConfig** commands.

Syntax

```
►►──CmsConfig verify──nodeName──┬─────────┬─────────────────────────────────►◄
                                └─cmsPort─┘
```

Parameters

nodeName
The client node name that is associated with the log files. For most client systems, only one node name is registered to the IBM Spectrum Protect server. However, on systems with multiple users, such as Linux client systems, there can be more than one client node name. This parameter is required.

cmsPort
The TCP/IP port number that is used to communicate with the client management service. Specify the port number if you did not use the default port number when you installed the client management service. The default port number is 9028. This parameter is optional.

Example for a Linux client system

Verify that the node definition for the node `SUSAN` is created correctly in the `client-configuration.xml` file.

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh verify SUSAN
```

During the verification process, you are prompted to enter the client node name or administrative user ID and password.

Output:

```
Verifying node.


Verifying the CMS service configuration for node SUSAN.
The CMS configuration looks correct.

Verifying the CMS service works correctly on port 9028.

Enter your user id: admin
Enter your password:

Connecting to CMS service and verifying resources.
The CMS service is working correctly.
Finished verifying node.
```

Example for a Windows client system

Verify that the node definition for the node `SUSAN` is created correctly in the `client-configuration.xml` file.

Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

Commands:

```
cmsconfig verify SUSAN
```

During the verification process, you are prompted to enter the client node name or administrative user ID and password.

Output:

```
Verifying node.


Verifying the CMS service configuration for node SUSAN.
The CMS configuration looks correct.

Verifying the CMS service works correctly on port 9028.

Enter your user id: admin
Enter your password:

Connecting to CMS service and verifying resources.
The CMS service is working correctly.
Finished verifying node.
```

`CmsConfig list` command:

Use the **`CmsConfig list`** command to show the client management service configuration.

Syntax

```
►►──CmsConfig list──────────────────────────────────────────────────────►◄
```

Example for a Linux client system

Show the configuration of the client management service. Then, view the output to ensure that you entered the command correctly.

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory.

Command:

```
./CmsConfig.sh list
```

Output:

```
Listing CMS configuration

server.example.com:1500 NO_SSL SUSAN
Capabilities: [LOG_QUERY]
    Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys

    Log File: /opt/tivoli/tsm/client/ba/bin/dsmerror.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252

    Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

Example for a Windows client system

Show the configuration of the client management service. Then, view the output to ensure that you entered the command correctly.

Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory.

Command:

```
cmsconfig list
```

Output:

```
Listing CMS configuration

server.example.com:1500 NO_SSL SUSAN
Capabilities: [LOG_QUERY]
    Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt

    Log File: C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252

    Log File: C:\Program Files\Tivoli\TSM\baclient\dsmsched.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

`CmsConfig help` command:

Use the **`CmsConfig help`** command to show the syntax of **`CmsConfig`** utility commands.

Syntax

```
►►──CmsConfig help──────────────────────────────────────────────────►◄
```

Example for a Linux client system

Issue the following command from the `/opt/tivoli/tsm/cms/bin` directory:

```
./CmsConfig help
```

Example for a Windows client system

Issue the following command from the `C:\Program Files\Tivoli\TSM\cms\bin` directory:

```
CmsConfig help
```

Advanced client management service capabilities:

By default, the IBM Spectrum Protect client management service collects information only from client log files. To initiate other client actions, you can access the Representational State Transfer (REST) API that is included with the client management service.

API developers can create REST applications to initiate the following client actions:

- Query and update client options files (for example, the `dsm.sys` file on Linux clients and the `dsm.opt` file on Linux and Windows clients).
- Query the status of the IBM Spectrum Protect client acceptor and the scheduler.
- Back up and restore files for a client node.
- Extend the capabilities of the client management service with scripts.

For detailed information about the client management service REST API, see the Client Management Services REST API Guide.

Chapter 13. Troubleshooting the Operations Center installation

If a problem occurs with the Operations Center installation and you cannot solve it, you can consult the descriptions of known problems for a possible solution.

Chinese, Japanese, or Korean fonts are displayed incorrectly

Chinese, Japanese, or Korean fonts are displayed incorrectly in the Operations Center on Red Hat Enterprise Linux 5.

Solution

Install the following font packages, which are available from Red Hat:
- fonts-chinese
- fonts-japanese
- fonts-korean

Chapter 14. Uninstalling the Operations Center

You can uninstall the Operations Center by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

Uninstalling the Operations Center by using a graphical wizard

You can uninstall the Operations Center by using the graphical wizard of IBM Installation Manager.

Procedure

1. Open IBM Installation Manager.

   In the directory where IBM Installation Manager is installed, go to the `eclipse` subdirectory (for example, `/opt/IBM/InstallationManager/eclipse`), and issue the following command:

   ```
   ./IBMIM
   ```
2. Click **Uninstall**.
3. Select the option for the Operations Center, and click **Next**.
4. Click **Uninstall**.
5. Click **Finish**.

Uninstalling the Operations Center in console mode

To uninstall the Operations Center by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:

   ```
   eclipse/tools
   ```

   For example:

   ```
   /opt/IBM/InstallationManager/eclipse/tools
   ```
2. From the `tools` directory, issue the following command:

   ```
   ./imcl -c
   ```
3. To uninstall, enter `5`.
4. Choose to uninstall from the IBM Spectrum Protect package group.
5. Enter `N` for Next.
6. Choose to uninstall the Operations Center package.
7. Enter `N` for Next.
8. Enter `U` for Uninstall.
9. Enter `F` for Finish.

Uninstalling the Operations Center in silent mode

To uninstall the Operations Center in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the Operations Center server. IBM Spectrum Protect includes a sample response file, `uninstall_response_sample.xml`, in the `input` directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

To uninstall the Operations Center, leave `modify="false"` set for the Operations Center entry in the response file.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
   ```
   eclipse/tools
   ```
   For example:
   ```
   /opt/IBM/InstallationManager/eclipse/tools
   ```
2. From the `tools` directory, issue the following command, where *response_file* represents the response file path, including the file name:
   ```
   ./imcl -input response_file -silent
   ```
   The following command is an example:
   ```
   ./imcl -input /tmp/input/uninstall_response.xml -silent
   ```

Chapter 15. Rolling back to a previous version of the Operations Center

By default, IBM Installation Manager saves earlier versions of a package to roll back to if you experience a problem with later versions of updates, fixes, or packages.

Before you begin

The rollback function is available only after the Operations Center is updated.

About this task

When IBM Installation Manager rolls back a package to a previous version, the current version of the package files is uninstalled, and an earlier version is reinstalled.

To roll back to a previous version, IBM Installation Manager must access files for that version. By default, these files are saved during each successive installation. Because the number of saved files increases with each installed version, you might want to delete these files from your system on a regular schedule. However, if you delete the files, you cannot roll back to a previous version.

To delete saved files or to update your preference for saving these files in future installations, complete the following steps:
1. In IBM Installation Manager, click **File** > **Preferences**.
2. On the Preferences page, click **Files for Rollback**, and specify your preference.

Procedure

To roll back to a previous version of the Operations Center, use the **Roll Back** function of IBM Installation Manager.

Part 3. Appendixes

Appendix A. Installation log files

If you experience errors during installation, these errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

You can view installation log files by clicking **File** > **View Log** from the Installation Manager tool. To collect these log files, click **Help** > **Export Data for Problem Analysis** from the Installation Manager tool.

Appendix B. 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 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Spectrum Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

SoftLayer® is a registered trademark of SoftLayer, Inc., an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability
These terms and conditions are in addition to any terms of use for the IBM website.

Personal use
You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use
You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled “Cookies, Web Beacons and Other Technologies,” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.

See the IBM Spectrum Protect glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

Product Number: 5725-W99
5725-W98
5725-X15

Printed in USA