

IBM Storage Protect
for Windows
8.1.20

Installation Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 191](#).

Edition Notice

This edition applies to version 8, release 1, modification 20 of IBM® Storage Protect (product numbers 5725-W98, 5725-W99, 5725-X15), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 1993, 2023.**

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

Contents

About this publication.....	vii
Who should read this guide.....	vii
Installable components.....	vii
Publications	viii
What's new.....	ix
Part 1. Installing and upgrading the server.....	1
Chapter 1. Planning to install the IBM Storage Protect server.....	3
What you should know first.....	3
What you should know about security before you install or upgrade the server.....	3
Applying security updates.....	7
Troubleshooting security updates.....	12
Planning for optimal performance.....	16
Planning server hardware and operating system.....	17
Planning server database disks.....	21
Planning server recovery log disks.....	23
Planning container storage pools.....	25
Planning for daily operations for directory-container storage pools.....	33
Planning DISK or FILE storage pools.....	35
Planning storage technology.....	38
Installation best practices.....	40
Minimum system requirements for the IBM Storage Protect server.....	42
IBM Installation Manager.....	42
Worksheets for planning details for the server.....	42
Capacity planning.....	43
Database space requirements.....	44
Recovery log space requirements.....	47
Monitoring space utilization for the database and recovery logs.....	58
Deleting installation rollback files	59
Server naming best practices.....	60
Chapter 2. Installing the server components.....	63
Obtaining the installation package.....	63
Using the installation wizard.....	64
Using the console installation wizard.....	65
Using silent mode.....	67
Installing server language packages.....	68
Server language locales.....	68
Configuring a language package.....	68
Updating a language package.....	69
Installing Open Snap Store Manager.....	69
Chapter 3. Taking the first steps after you install IBM Storage Protect.....	71
Creating the user ID and directories for the server instance.....	71
Configuring the IBM Storage Protect server.....	73
Using the configuration wizard.....	73
Using the manual configuration steps.....	75
Configuring server options for server database maintenance.....	79

Starting the server instance on Windows systems.....	80
Configuring the server to start as a Windows service.....	81
Starting the server as a Windows service.....	82
Manually creating and configuring a Windows service.....	82
Starting the server in the foreground.....	83
Starting the server in maintenance mode.....	83
Stopping the server.....	84
Registering licenses.....	84
Preparing the server for database backup operations	85
Running multiple server instances on a single system.....	85
Monitoring the server.....	86
 Chapter 4. Installing an IBM Storage Protect fix pack.....	89
Applying a fix pack to V8 for Windows in a clustered environment.....	91
 Chapter 5. Upgrading the server to V8.1.....	95
Upgrading to V8.1.....	95
Planning the upgrade.....	96
Preparing the system.....	96
Installing the server and verifying the upgrade.....	99
Upgrading the server in a clustered environment.....	102
Upgrading a V7.1 server to V8.1 in a clustered environment	102
Upgrading IBM Storage Protect servers in a clustered HADR environment.....	105
 Chapter 6. Reverting to the previous server version in a cluster configuration.....	107
Steps for reverting to the previous server version.....	107
 Chapter 7. Reference: Db2 commands for server databases.....	109
 Chapter 8. Uninstalling IBM Storage Protect.....	113
Uninstalling IBM Storage Protect by using a graphical wizard.....	113
Uninstalling IBM Storage Protect in console mode.....	113
Uninstalling IBM Storage Protect in silent mode.....	114
Uninstalling and reinstalling IBM Storage Protect.....	114
Uninstalling IBM Installation Manager.....	115
 Part 2. Installing and upgrading the Operations Center.....	117
 Chapter 9. Planning to install the Operations Center.....	119
System requirements for the Operations Center.....	119
Operations Center computer requirements.....	120
Hub and spoke server requirements.....	120
Operating system requirements.....	122
Web browser requirements.....	123
Language requirements.....	123
Requirements and limitations for IBM Storage Protect client management services.....	123
Administrator IDs that the Operations Center requires.....	125
IBM Installation Manager.....	125
Installation checklist.....	126
 Chapter 10. Installing the Operations Center.....	129
Obtaining the Operations Center installation package.....	129
Installing the Operations Center by using a graphical wizard.....	129
Installing the Operations Center in console mode.....	130
Installing the Operations Center in silent mode.....	130
Encrypting passwords in silent installation response files.....	131

Chapter 11. Upgrading the Operations Center.....	133
Chapter 12. Getting started with the Operations Center.....	135
Configuring the Operations Center.....	135
Designating the hub server.....	136
Adding a spoke server.....	136
Sending email alerts to administrators.....	137
Adding customized text to the login screen.....	140
Configuring the Operations Center web server to use the standard TCP/IP secure port.....	140
Enabling REST services.....	141
Configuring for secure communication.....	141
Between the Operations Center and the hub server by using self-signed certificates.....	142
Between the Operations Center and the hub server by using CA-signed certificates.....	144
Between the hub server and a spoke server.....	145
Between the Operations Center and web browsers.....	147
Deleting and reassigning the password for the Operations Center truststore file.....	158
Starting and stopping the web server.....	160
Opening the Operations Center.....	160
Collecting diagnostic information with the client management service.....	161
Collecting diagnostic information on 8.1.13 or later versions	161
Collecting diagnostic information on versions earlier than 8.1.13.....	162
Chapter 13. Uninstalling the Operations Center.....	183
Uninstalling the Operations Center by using a graphical wizard.....	183
Uninstalling the Operations Center in console mode.....	183
Uninstalling the Operations Center in silent mode.....	183
Chapter 14. Rolling back to a previous version of the Operations Center.....	185
Appendix A. Installation log files.....	187
Appendix B. Accessibility.....	189
Notices.....	191
Glossary.....	195
Index.....	197

About this publication

This publication contains installation and configuration instructions for the IBM Storage 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 Storage Protect server or Operations Center.

Installable components

The IBM Storage Protect server and licenses are required components.

These components are in several different installation packages.

Table 1. IBM Storage Protect installable components		
IBM Storage 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.	“Installing IBM Storage Protect by using the installation wizard” on page 64
Language package (optional)	Each language package (one for each language) contains language-specific information for the server.	See “Installing server language packages” on page 68 .
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 .

Table 1. IBM Storage Protect installable components (continued)

IBM Storage Protect component	Description	Additional information
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 Storage 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) .
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 117.

Publications

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

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

What's new in this release

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

For a list of new features and updates, see [What's new](#).

If changes were made in the documentation, they are indicated by a vertical bar (|) in the margin.

Part 1. Installing and upgrading the server

Install and upgrade the IBM Storage 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 Storage Protect server-managed storage.

What you should know first

Before installing IBM Storage 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 cannot install and run the server on a system that already has Db2® installed on it, whether Db2 was installed by itself or as part of some other application. The server requires the installation and use of the Db2 version that is packaged with the server. No other version of Db2 can exist on the system.

You can install the IBM Storage Protect server on a domain controller. The server can have heavy processor usage, however, and that might affect and stall other applications.

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 Storage Protect, or alter the Db2 environment for IBM Storage Protect in other ways, such as with other products. The 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 Storage 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

Review information about the enhanced security features in the IBM Storage Protect server and the requirements for updating your environment.

Before you begin

Beginning in Version 8.1.2, enhancements were added to IBM Storage Protect that enforce stricter security settings. Before you install or upgrade IBM Storage Protect, complete the following steps:

- In IBM Documentation, in the *What's New* topic, review the information in the Security sections to learn about security updates for each version.
- If you have previous versions of the server in your environment, review the restrictions and known issues in [technote 562939](#). To avoid these restrictions and take advantage of the latest security enhancements, plan to update all IBM Storage Protect servers and backup-archive clients in your environment to the latest version.
- Verify that you backed up the following directories and files, which are required to restore the server:
 - Server options file (`dsmserv.opt`)
 - Device configuration file (for example, `devconf.dat`)
 - Volume history file (for example, `volhist.dat`)
 - Master encryption key files (`dsmkeydb.kdb` or `dsmkeydb.sth`)
 - Server certificate and private key files (`cert.kbd` or `cert.sth`)

Security enhancements

The following security enhancements were added beginning in V8.1.2:

Security protocol that uses Transport Layer Security (TLS)

IBM Storage Protect V8.1.2 and later software has an improved security protocol that uses TLS Version 1.2 or later for authentication between the server, storage agent, and backup-archive clients.

Beginning with IBM Storage Protect V8.1.11, you can enable the TLS 1.3 protocol to secure communications between servers, clients, and storage agents. To use TLS 1.3, both parties in the communication session must use TLS 1.3. If either party uses TLS 1.2, then both parties use TLS 1.2 by default.

Automatic Secure Sockets Layer (SSL) configuration and distribution of certificates

Servers, storage agents, and clients using V8.1.2 or later software are automatically configured to authenticate with each other by using TLS.

Using the new protocol, each server, storage agent, and client has a unique self-signed certificate that is used to authenticate and allow TLS connections. IBM Storage Protect self-signed certificates enable secure authentication between entities, enable strong encryption for data transmission, and automatically distribute public keys to client nodes. Certificates are automatically exchanged between all clients, storage agents, and servers that use V8.1.2 or later software. You do not have to manually configure TLS or manually install the certificates for every client. The new TLS enhancements do not require options changes, and certificates are transferred to clients automatically upon first connection unless you are using a single administrator ID to access multiple systems.

By default, self-signed certificates are distributed, but you can optionally use other configurations such as certificates that are signed by a certificate authority. For more information about using certificates, see *SSL and TLS communication* in IBM Documentation.

Combination of TCP/IP and TLS protocols for secure communication and minimal impact to performance

In previous versions of IBM Storage Protect software, you had to choose either TLS or TCP/IP to encrypt all communication. The new security protocol uses a combination of TCP/IP and TLS to secure communication between servers, clients, and storage agents. By default, TLS is used only to encrypt authentication and metadata, while TCP/IP is used for data transmission. Since TLS encryption is primarily used for authentication only, performance for backup and restore operations is not affected.

Optionally, you can use TLS to encrypt data transmission by using the **SSL** client option for client-to-server communication, and the **SSL** parameter in the **UPDATE SERVER** command for server-to-server communication.

Backward compatibility makes it easier to plan upgrades in batches

Upgraded versions of IBM Storage Protect servers and clients can continue to connect to older versions when the **SESSIONSECURITY** parameter is set to **TRANSITIONAL**.

You are not required to update backup-archive clients to V8.1.2 or later before you upgrade servers. After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the **TRANSITIONAL** value until the entity meets the requirements for the **STRICT** value. Similarly, you can upgrade backup-archive clients to V8.1.2 or later before you upgrade your IBM Storage Protect servers, but you are not required to upgrade servers first. Communication between servers and clients that are using different versions is not interrupted. However, you will not have the benefits of the security enhancements until both clients and servers are upgraded.

Enforce strict security with the **SESSIONSECURITY** parameter

To use the new security protocol, the server, client node, or administrator entities must be using IBM Storage Protect software that supports the **SESSIONSECURITY** parameter. Session security is the level of security that is used for communication among IBM Storage Protect client nodes, administrative clients, and servers. You can specify the following values for this parameter:

STRICT

Enforces the highest level of security for communication between IBM Storage Protect servers, nodes, and administrators, which is currently TLS 1.2.

TRANSITIONAL

Specifies that the existing communication protocol (for example, TCP/IP) is used until you update your IBM Storage Protect software to V8.1.2 or later. This is the default. When **SESSIONSECURITY=TRANSITIONAL**, stricter security settings are automatically enforced as higher versions of the TLS protocol are used and as the software is updated to V8.1.2 or later. After a node, administrator, or server meets the requirements for the STRICT value, session security is automatically updated to the STRICT value, and the entity can no longer authenticate by using a previous version of the client or earlier TLS protocols.

If **SESSIONSECURITY=TRANSITIONAL** and the server, node, or administrator has never met the requirements for the STRICT value, the server, node, or administrator will continue to authenticate by using the TRANSITIONAL value. However, after the server, node, or administrator meets the requirements for the STRICT value, the **SESSIONSECURITY** parameter value automatically updates from TRANSITIONAL to STRICT. Then, the server, node, or administrator can no longer authenticate by using a version of the client or an SSL/TLS protocol that does not meet the requirements for STRICT.

Restriction: After an administrator successfully authenticates with a server by using IBM Storage Protect V8.1.2 or later software or Tivoli® Storage Manager V7.1.8 or later software, the administrator can no longer authenticate with the same server by using client or server versions earlier than V8.1.2 or V7.1.8. This restriction also applies to the destination server when you use functions such as command routing, server-to-server export that authenticates with the destination IBM Storage Protect server as an administrator from another server, administrator connections using the Operations Center, and connections from the administrative command-line client.

For client and administrative sessions, administrative command routing sessions might fail unless the administrator ID has already acquired certificates for all servers to which the administrator ID will connect. Administrators that authenticate by using the **dsmadmc** command, **dsmc** command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:

- Ensure that all IBM Storage Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system.
- If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.

Before you upgrade

Before you upgrade a server, review the guidelines in the following checklist.

Table 2. Planning checklist

Guideline	Description
<p>Back up the following server files:</p> <ul style="list-style-type: none"> • Key databases (cert.kdb and dsmkeydb.kdb) • Stash files (cert.sth and dsmkeydb.sth) 	<p>Beginning with IBM Storage Protect Version 8.1.2, a master encryption key is automatically generated when you start the server if the master encryption key did not previously exist.</p> <p>The master encryption key is stored in a key database, dsmkeydb.kdb. Server certificates are still stored in the cert.kdb key database and accessed by the stash file cert.sth. You must protect both the key databases (cert.kdb and dsmkeydb.kdb) and the stash files (cert.sth and dsmkeydb.sth) that provide access to each of the key databases. By default, the BACKUP DB command protects the master encryption key in the same manner in which the volume history and devconfig files are protected. You must remember the database backup password to restore the database. The IBM Storage Protect server dsmseiv . pwd file, which was used to store the master encryption key in previous releases, is no longer used.</p>
<p>Carefully plan upgrades for administrator IDs</p>	<p>Identify all systems that administrator accounts use to log in for administration purposes.</p> <p>After a successful authentication to V8.1.2 or later software, administrators cannot authenticate to earlier versions of IBM Storage Protect software on the same server. If a single administrator ID is used to log in to multiple systems, plan to upgrade all of those systems with V8.1.2 or later software to ensure that the certificate is installed on all systems that the administrator logs in to.</p> <p>Tip: You will not get locked out of a server if the SESSIONSECURITY parameter for all of your administrator IDs is updated to the STRICT value. You can manually import the server's public certificate to a client from which you issue the dsmadm command.</p>

Table 2. Planning checklist (continued)

Guideline	Description
<p>If you're using TLS with previous versions of the client that use the "TSM Server SelfSigned Key" (cert.arm) certificate, update your clients to V8.1.4 or later.</p>	<p>In releases prior to V7.1.8, the default certificate was labeled "TSM Server SelfSigned Key" and had an MD5 signature, which does not support the TLS 1.2 or later protocol that is required by default for V8.1.2 or later clients and the Operations Center. To resolve this issue, complete one of the following steps:</p> <ul style="list-style-type: none"> • Upgrade the server to V8.1.4 or later. Beginning with V8.1.4, servers that use the MD5-signed certificate as the default are automatically updated to use a default certificate with a SHA signature that is labeled TSM Server SelfSigned SHA Key. A copy of the new default certificate is stored in the cert256.arm file, which is located in the server instance directory. <p>Tip: Before you update the server to use the new default certificate with a SHA signature, distribute the cert256.arm file to clients to prevent client backup failures. Each client must obtain and import the new certificate before they can connect to a server that is using the new default SHA certificate. You do not need to remove previous certificates.</p> <ul style="list-style-type: none"> • To manually update your default certificate, follow the instructions in technote 562939.

What to do next

- Follow the procedure in [“Applying security updates” on page 7](#) to install or upgrade an IBM Storage Protect server.
- For information about troubleshooting communication issues related to security updates, see [“Troubleshooting security updates” on page 12](#).
- For FAQ information, see [FAQ - Security updates in IBM Storage Protect](#).
- For information about using the IBM Storage Protect backup-archive web client in the new security environment, see [technote 728037](#).

Applying security updates

Apply security updates that are delivered with new releases of IBM Storage Protect.

Before you begin

Review the following information:

- For details about security updates delivered with a release, see the *What's New* topic in IBM Documentation.
- For information about the updates and any restrictions that can apply, see [“What you should know about security before you install or upgrade the server” on page 3](#).
- To determine the order in which you upgrade the servers and clients in your environment, answer the following questions:

Table 3. Questions for consideration before upgrading	
Question	Consideration
What is the role of the server in the configuration?	In general, you can upgrade the IBM Storage Protect servers in your environment first and then upgrade backup-archive clients. However, in certain circumstances, for example, if you use command routing functions, the server can act as the client in your configuration. In that instance, to prevent communication issues, the suggested approach is to upgrade clients first. For information about different scenarios, see Upgrade scenarios .
What systems are used for administrator authentication?	<p>For administrator accounts, the sequence in which you upgrade is important to prevent authentication issues.</p> <ul style="list-style-type: none"> – Clients on multiple systems that log on by using the same ID (either node or administrative ID) must be upgraded at the same time. Server certificates are transferred to clients automatically upon first connection. – Before you upgrade your server, consider all endpoints that the administrator uses to connect to for administration purposes. If a single administrative ID is used to access multiple systems, ensure that the server's certificate is installed on each system. – After an administrator ID authenticates successfully with the server by using IBM Storage Protect V8.1.2 or later software or Tivoli Storage Manager V7.1.8 or later software, the administrator can no longer authenticate with that server by using client or server versions earlier than V8.1.2 or V7.1.8. This is also true for a destination server when you authenticate with that destination IBM Storage Protect server as an administrator from another server. For example, this is true when you use the following functions: <ul style="list-style-type: none"> - Command routing - Server-to-server export - Connecting from an administrative client in the Operations Center

Table 3. Questions for consideration before upgrading (continued)

Question	Consideration
In what sequence should I upgrade my systems?	<ul style="list-style-type: none"> – If you upgrade servers before you upgrade client nodes: <ul style="list-style-type: none"> - Upgrade the hub server first and then any spoke servers. - When you upgrade a server to V8.1.2 or later, nodes and administrators that use earlier versions of the software can continue to communicate with the new server by using the existing communication protocol. The SESSIONSECURITY is set to TRANSITIONAL and if the server, node, or administrator has never met the requirements for the STRICT value, the server, node, or administrator continues to authenticate by using the TRANSITIONAL value. However, as soon as the server, node, or administrator meets the requirements for the STRICT value, the SESSIONSECURITY parameter value automatically updates from TRANSITIONAL to STRICT. – If you upgrade client nodes before you upgrade servers: <ul style="list-style-type: none"> - Upgrade administrative clients first, and then upgrade non-administrative clients. Clients at later release levels continue to communicate with servers at earlier levels. Important: If you upgrade any one of the administrative clients in your environment, all other clients that use the same ID as the upgraded client must be upgraded at the same time. - It is not necessary to upgrade all of your non-administrative clients at the same time, unless multiple clients are using the same ID to log on. Then, all other clients that use the same ID as the upgraded client must be upgraded at the same time and the server's certificate must be installed on each system.

About this task

If your environment includes IBM Storage Protect backup-archive clients or IBM Storage Protect servers that are earlier than V7.1.8 or V8.1.2, you might have to customize your configuration to ensure that communication between servers and clients is not interrupted. Follow the default procedure in this topic for installing or upgrading your environment.

Review [Upgrade scenarios](#) for other example scenarios that might apply to your environment.

Tip: To take advantage of the latest security enhancements, plan to update all IBM Storage Protect servers and backup-archive clients in your environment to the latest release level.

Procedure

1. Install or upgrade IBM Storage Protect servers in your environment. For more information, see the *Installing and upgrading the server* topic in IBM Documentation.
 - a) Upgrade the Operations Center and the hub server. For more information, see [Part 2, “Installing and upgrading the Operations Center,”](#) on page 117.
 - b) Upgrade spoke servers.
 - c) Configure or verify server-to-server communications. For more information, see the following topics:
 - The *UPDATE SERVER* command in IBM Documentation.
 - The *Configuring SSL communications between the hub server and a spoke server* topic in IBM Documentation.
 - The *Configuring the server to connect to another server by using SSL* topic in IBM Documentation.

Tip:

- Beginning in IBM Storage Protect V8.1.2 and Tivoli Storage Manager V7.1.8, the **SSL** parameter uses SSL to encrypt communication with the specified server even if the **SSL** parameter is set to NO.
 - Beginning with V8.1.4, certificates are automatically configured between storage agents, library clients, and library manager servers. Certificates are exchanged the first time a server-to-server connection is established to a server with enhanced security.
2. Install or upgrade administrative clients. For more information, see the *Installing and configuring clients* topic in IBM Documentation.
 3. Enable secure communications between all systems that administrators use to log in for administration purposes.
 - Ensure that the IBM Storage Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later.
 - If an administrative ID logs on from multiple systems, ensure that the server's certificate is installed on each system.
 4. Install or upgrade non-administrative clients. For more information, see the *Installing and configuring clients* topic in IBM Documentation.

Remember: You can upgrade your non-administrative clients in phases. You can continue to connect to servers at later release levels from clients at earlier release levels by issuing the **UPDATE NODE** command and setting the **SESSIONSECURITY** parameter to TRANSITIONAL for each node.

```
update node nodename sessionsecurity=transitional
```

What to do next

Other upgrade scenarios might apply to your environment. Review example upgrade scenarios in the following table.

Table 4. Upgrade scenarios		
Scenario	Considerations	Suggested upgrade approach
I use administrative command routing functions to route commands to one or more servers. I want to connect to an IBM Storage Protect server that is earlier than V8.1.2.	<ul style="list-style-type: none"> With command routing, the server can act as the administrative client. Command routing uses the ID and the password of the administrator who is issuing the command. If you use a single administrative ID to access multiple systems, ensure that the server's certificate is installed on each system. 	<ul style="list-style-type: none"> Upgrade the administrative client first. <ul style="list-style-type: none"> Important: Clients on multiple systems that log on by using the same node or administrative ID must be upgraded at the same time. On each server to which commands are being routed, verify that the following information is configured: <ul style="list-style-type: none"> The same administrator ID and password The required administrative authority on each server The required certificates are installed Upgrade the servers that the administrator account uses to log on to V8.1.2 or later.
My administrative client is at the latest release version, and I use the same administrator ID to authenticate to different systems by using the dsmadmc command. I have authenticated successfully to an IBM Storage Protect server in my environment that is running at the latest version. I now want to authenticate to a server at a version earlier than V8.1.2.	<ul style="list-style-type: none"> After an administrator authenticates to an IBM Storage Protect server V8.1.2 or later by using a version of the client at V8.1.2 or later, the administrative ID can only authenticate with that server on clients or servers that are using V8.1.2 or later. If you use a single administrative ID to access multiple systems, plan to upgrade all of those systems with V8.1.2 or later software to ensure that the server's certificate is installed on all systems to which the administrator logs on. 	<ul style="list-style-type: none"> Ensure that all IBM Storage Protect software that the administrators use to log on is upgraded to V8.1.2 or later. The preferred action is to upgrade all the servers in your environment to the latest version. If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.
The IBM Storage Protect server is already upgraded to the latest release level. I have an administrative client at release level V8.1.0 and I want to connect to the server from the Operations Center.	<ul style="list-style-type: none"> If you upgrade any one of the administrative clients in your environment, all other clients that use the same ID as the upgraded client must be upgraded at the same time. 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, authority level, and required certificates. 	<ul style="list-style-type: none"> On each server, verify that the following information is set up: <ul style="list-style-type: none"> The same administrator ID and password The required administrative authority on each server The required certificates Upgrade non-administrative clients in a phased manner.

Table 4. Upgrade scenarios (continued)

Scenario	Considerations	Suggested upgrade approach
I use node replication to protect my data.	<ul style="list-style-type: none"> The replication heartbeat initiates a certificate exchange when the first server-to-server connection is established after you upgrade the server. 	<ul style="list-style-type: none"> Upgrade your servers before you upgrade your clients; follow the default procedure.
I want to upgrade my backup-archive clients before I upgrade my servers.	<ul style="list-style-type: none"> After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the TRANSITIONAL value until the entity meets the requirements for the STRICT value. Communication between servers and clients will not be interrupted. 	<ul style="list-style-type: none"> If you upgrade your clients before you upgrade your servers, upgrade administrative clients first, and then upgrade non-administrative clients. Clients at later release levels continue to communicate with servers at earlier levels.

Troubleshooting security updates

Troubleshoot issues that might occur after you upgrade IBM Storage Protect.

Symptom	Resolution
An administrator account cannot log in to a system that is using software earlier than V8.1.2.	<p>After an administrator successfully authenticates with the server by using IBM Storage Protect V8.1.2 or later software, the administrator can no longer authenticate with that server that uses client or server versions earlier than V8.1.2. This restriction also applies to the destination server when you use functions such as command routing, server-to-server export that authenticates with the destination IBM Storage Protect server as an administrator from another server, administrator connections that use the Operations Center, and connections from the administrative command-line client.</p> <p>To resolve authentication issues for administrators, complete the following steps:</p> <ol style="list-style-type: none"> 1. Identify all systems from which administrators log in and which use the administrative ID to log in. Upgrade the system software to IBM Storage Protect V8.1.2 or later, and ensure that the server's certificate is installed on each system. 2. Set the administrator's SESSIONSECURITY parameter value to TRANSITIONAL by issuing the command <code>update admin admin_name sessionsecurity=transitional</code> 3. Retry the administrator connection. <p>Tip: If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.</p>
Certificate distribution failed for a node, administrator, or server.	A node, administrator, or server that is using V8.1.2 or later software has a SESSIONSECURITY value of STRICT, but you have to reset the value to TRANSITIONAL to retry certificate distribution.

Symptom	Resolution
	<p>When using the new protocol, the automatic transfer of a server's public certificate is performed only on the first connection to a server with enhanced security. After the first connection, the SESSIONSECURITY parameter value of a node changes from TRANSITIONAL to STRICT. You can temporarily update a node, administrator, or server to TRANSITIONAL to allow another automatic transfer of the certificate. While in TRANSITIONAL, the next connection automatically transfers the certificate if needed and resets the SESSIONSECURITY parameter to STRICT.</p> <p>Update the value of the SESSIONSECURITY parameter to TRANSITIONAL by issuing one of the following commands:</p> <ul style="list-style-type: none"> • For client nodes, issue: <code>update node <i>node_name</i> sessionsecurity=transitional</code> • For administrators, issue: <code>update admin <i>admin_name</i> sessionsecurity=transitional</code> • For servers, issue: <code>update server <i>server_name</i> sessionsecurity=transitional</code> <p>Alternatively, you can manually transfer and import the public certificate by using the <code>dsmdcert</code> utility to issue the following commands:</p> <pre>openssl s_client -connect tapsrv04:1500 -showcerts > tapsrv04.arm</pre> <pre>dsmdcert -add -server tapsrv04 -file tapsrv04.arm</pre> <p>If you are using CA-signed certificates, you must install the CA-root and any CA-intermediate certificates on each key database for the client, server, and storage agent that initiates SSL communication.</p>
Certificate exchange between IBM Storage Protect servers was not successful.	<p>When using the new protocol, the automatic transfer of a server's public certificate is performed only on the first connection to a server with enhanced security. After the first connection, the SESSIONSECURITY parameter value of a server changes from TRANSITIONAL to STRICT. Retry certificate exchange between two IBM Storage Protect servers. For information, see <i>Retrying certificate exchange between servers</i>.</p>
Certificate exchange between an IBM Storage Protect server and a client node was not successful.	<p>When using the new protocol, the automatic transfer of a server's public certificate is performed only on the first connection to a server with enhanced security. After the first connection, the SESSIONSECURITY parameter value of a node changes from TRANSITIONAL to STRICT. To retry certificate exchange between clients and servers at versions earlier than V8.1.2, complete these steps:</p> <ol style="list-style-type: none"> 1. For existing clients that are configured to use SSL with the <code>cert.arm</code> certificate, reconfigure them to use the <code>cert256.arm</code> certificate. For instructions, see <i>Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL</i> in IBM Documentation. 2. Update the default certificate by issuing the following command from the server instance directory: <pre>gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed -label "TSM Server SelfSigned SHA Key"</pre> 3. Restart the server.

Symptom	Resolution
	<p>For clients and servers at V8.1.2 and later, the certificates are automatically distributed. If communication between clients or servers fails, complete these steps to retry certificate acquisition:</p> <ol style="list-style-type: none"> For nodes and administrators, set the SESSIONSECURITY parameter to TRANSITIONAL by issuing the following commands for each node or administrator that you want to retry: <pre>update node nodename sessionsecurity=transitional update admin adminname sessionsecurity=transitional</pre> <p>Tip: Administrators that authenticate by using the dsmadm command, dsmc command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:</p> <ul style="list-style-type: none"> Ensure that all IBM Storage Protect software that the administrator account uses to log in is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system before the administrator account is used for command routing. After an administrator authenticates to a V8.1.2 or later server by using a V8.1.2 or later client, the administrator can authenticate only on clients or servers that are using V8.1.2 or later. An administrator command can be issued from any system. If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software. For storage agents, update the STASESSIONSECURITY option in the storage agent options file <code>dsmsta.opt</code> by changing the STRICT value to TRANSITIONAL. Restart the servers. Certificate changes do not take effect until you restart the servers or storage agents. If you are still unable to exchange certificates after completing Steps 1-4, manually add the certificates to the servers and storage agents and restart them. For instructions, see <i>Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL</i> in IBM Documentation.
You want to manually distribute certificates to client systems.	<p>The IBM Storage Protect server administrator can automatically deploy a backup-archive client to update workstations where the backup-archive client is already installed. For information, see <i>Automatic backup-archive client deployment</i> in IBM Documentation.</p> <p>To manually add certificates to clients, see <i>Configuring IBM Storage Protect client/server communication with Secure Sockets Layer</i> in IBM Documentation.</p>
You want to reset certificates for client-to-client sessions.	<p>The <code>dsmcert</code> utility that is installed with the IBM Storage Protect backup-archive client is used to create a certificate store for server certificates. Use the <code>dsmcert</code> utility to delete the files and re-import the certificates.</p>
As a root user, you want to allow non-root users to manage your files.	<p>The trusted communications agent (TCA), previously used by non-root users in V8.1.0 and V7.1.6 and earlier IBM Storage Protect clients, is no longer available. Root users can use the following methods to allow non-root users to manage their files:</p>

Symptom	Resolution
	<p>Help desk method</p> <p>With the help desk method, the root user runs all backup and restore operations. The non-root user must contact the root user to request certain files to be backed up or restored.</p> <p>Authorized user method</p> <p>With the authorized user method, a non-root user is given read/write access to the password store by using the <code>passworddir</code> option to point to a password location that is readable and writable by the non-root user. This method allows non-root users to back up and restore their own files, use encryption, and manage their passwords with the <code>passwordaccess generate</code> option.</p> <p>For more information, see <i>Enable non-root users to manage their own data</i> in IBM Documentation.</p> <p>If neither of these methods are satisfactory, you must use the earlier clients that included the TCA.</p>
You want to resolve GSKit compatibility issues.	<p>When multiple applications that use GSKit are installed on the same system, incompatibility issues might occur. To resolve these issues, see the following information:</p> <ul style="list-style-type: none"> • For IBM Storage Protect clients, see Technote 2011742. • For Db2, see Technote 7050721. • For IBM Storage Protect server, see Technote 2007298. • For IBM Storage Protect server and client on the same Windows system, see Technote 7050721.

For more information about troubleshooting security updates, see [technote 2004844](#).

Retrying certificate exchange between servers

If the certificate exchange between servers fails, you can attempt another exchange.

Procedure

1. Remove the certificate from the partner server's database by issuing the following command on both servers:

```
update server servername forcesync=yes
```

Tip: The server might be using the wrong certificate if you are still getting error messages for each server-to-server session after you have completed the steps in this task and restarted the servers. If you determine that the server is attempting to use the wrong certificate, delete the certificate from the key database by issuing the following command:

```
gsk8capicmd_64 -cert -delete -db cert.kdb -stashed -label certificate_labelname
```

2. Delete the server definition by issuing the **DELETE SERVER** command for both the server and the partner server. If you cannot delete the server definition, you must configure the certificates manually. For instructions about manually configuring certificates, see *Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL* in IBM Documentation.
3. To reacquire the certificate, cross-define the servers to each other and allow them to exchange certificates by issuing the following commands on both servers:

```
set crossdefine on
set serverhladdress hladdress
```

```
set serverlladdress lladdress
set serverpassword password
```

4. Issue the following command on one of the servers that you are cross defining:

```
define server servername crossdefine=yes ssl=yes
```

5. Repeat step 3 for all other Version 8.1.2 or later server pairs.
6. Restart the servers.
7. To verify that certificates were exchanged, issue the following command from the server instance directory of each server that you want to verify:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

Example output:

```
example.website.com:1542:0
```

Tip: If you use replication, the replication heartbeat runs approximately every 5 minutes and initiates a certificate exchange during the first connection after you upgrade the server. This connection causes messages ANR8583E and ANR8599W to appear in the log once, before a certificate exchange takes place. If you do not use replication, certificates are exchanged the first time a server-to-server session is initiated, except for server configurations without a server defined on both computers.

8. For servers that are defined as a virtual volume, complete the following steps:
 - a) Remove the partner certificate from the server's database by issuing the following command on both servers:

```
update server servername forcesync=yes
```

- b) Ensure that the same password is used for the server password value on the **DEFINE SERVER** command on the source server, the password value on the **REGISTER NODE** command on the virtual volume server, and the **SET SERVERPASSWORD** value on the virtual volume server. If necessary, update a password by using the **UPDATE SERVER**, **UPDATE NODE**, or **SET SERVERPASSWORD** commands, respectively. Certificates are exchanged after the first client backup operation from the virtual volume server to the source server.
9. If you are still unable to exchange certificates between servers, complete the following steps:
 - a) In the server definition for each of the communicating servers, verify that you specified a server name that matches the name that was set by issuing the **SET SERVERNAME** command on the partner server.
 - b) Verify that server definitions have passwords that are specified with the **SET SERVERPASSWORD** command. The passwords must match the value that is specified with the **SET SERVERNAME** command for the partner server.
 - c) After completing steps a and b, reissue the following command:

```
update server servername forcesync=yes
```
 - d) Retry steps 1 through 3.

Planning for optimal performance

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

About this task

The optimal IBM Storage Protect environment is set up by using the [IBM Storage Protect Blueprints](#).

Procedure

1. Review [“What you should know first”](#) on page 3.
2. Review each of the following subsections.

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
<p>Does the operating system and hardware meet or exceed requirements?</p> <ul style="list-style-type: none"> • Number and speed of processors • System memory • Supported operating system level 	<p>If you are using the minimum required amount of memory, you can support a minimal workload.</p> <p>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.</p> <p>If you run multiple servers on the system, add the requirements for each server to get the requirements for the system.</p>	<p>Review operating system requirements at technote 84861.</p> <p>Additionally, review the guidance in Tuning tasks for operating systems and other applications.</p> <p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For more information about sizing requirements for the server and storage, see the IBM Storage Protect Blueprint.</p>
<p>Are disks configured for optimal performance?</p>	<p>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.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • "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?	<p>Heavier workloads and advanced features such as data deduplication and data replication require more than the minimum system memory that is specified in the system requirements document.</p> <p>For databases that are not enabled for data deduplication, use the following guidelines to specify memory requirements:</p> <ul style="list-style-type: none"> • For daily ingest of data of less than 1 TB, you need 24 GB of memory. • For daily ingest of data of 1 TB - 10 TB, you need 64 GB of memory. • For daily ingest of data of 10 TB - 30 TB, you need 192 GB of memory. • For daily ingest of data up to 100 TB, you need 384 GB of memory. <p>Ensure that you allocate extra space for the active log and the archive log for replication processing.</p>	<p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication • Memory requirements

Question	Tasks, characteristics, options, or settings	More information
Does the system have enough host bus adapters (HBAs) to handle the data operations that the IBM Storage Protect server must run simultaneously?	<p>Understand what operations require use of HBAs at the same time.</p> <p>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.</p>	See Tuning HBA capacity .
Is network bandwidth greater than the planned maximum throughput for backups?	<p>Network bandwidth must allow the system to complete operations such as backups in the time that is allowed or that meets service level commitments.</p> <p>For data replication, network bandwidth must be greater than the planned maximum throughput.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Tuning network performance • Checklist for node replication
Are you using a preferred file system for IBM Storage 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. For more information about the preferred file system for your operating system, see IBM Storage Protect server-supported file systems .	For more information, see Configuring the operating system for disk performance .

Question	Tasks, characteristics, options, or settings	More information
<p>Are you planning to configure enough paging space?</p>	<p>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.</p> <p>Restriction: Do not use paging space to add memory to your system. Paging space is intended to provide only a limited and temporary extension of space. If your system uses paging space, system memory is full and must be extended.</p> <p>Paging space is automatically configured.</p>	

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?	<p>Do not use the following drives for the IBM Storage Protect database:</p> <ul style="list-style-type: none"> • Nearline SAS (NL-SAS) • Serial Advanced Technology Attachment (SATA) • Parallel Advanced Technology Attachment (PATA) <p>Do not use internal disks that are included by default in most server hardware.</p> <p>Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.</p> <p>If you plan to use the data deduplication functions of IBM Storage Protect, focus on disk performance in terms of I/O operations per second (IOPS).</p>	For more information, see Checklist for data deduplication .
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?	<p>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.</p> <p>Tip: The database and the active log can share an array when you use solid-state drive (SSD) technology.</p>	
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?	<p>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.</p> <p>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.</p>	

Question	Tasks, characteristics, options, or settings	More information
<p>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?</p>	<p>If you can set the strip size or segment size, use 64 KB or 128 KB sizes on disk systems for the database.</p>	<p>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.</p>
<p>Are you planning to create at least four directories, also called storage paths, on four separate LUNs for the database?</p> <p>Create one directory per distinct array on the subsystem. If you have fewer than three arrays, create a separate LUN volume within the array.</p>	<p>Heavier workloads and use of some features require more database storage paths than the minimum requirements.</p> <p>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.</p> <p>Use the following guidelines to create directories in the server database:</p> <ul style="list-style-type: none"> • For server databases less than 2 TB, you need 4 directories. • For server databases with a size of 2 - 4 TB, you need 8 directories. • For server databases greater than 4 TB, you need 12 directories. <p>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.</p> <p>Use the <i>DB2_PARALLEL_IO</i> 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 <i>DB2_PARALLEL_IO</i> 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.</p>	<p>For more information, see the following topics:</p> <p>For help with forecasting growth when the server deduplicates data, see technote 1596944.</p> <p>For the most recent information about database size, database reorganization, and performance considerations for IBM Storage Protect servers, see technote 1683633.</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM Db2 registry variables.</p>

Question	Tasks, characteristics, options, or settings	More information
Are all directories for the database the same size?	<p>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.</p> <p>This guideline also applies if you must add storage paths after the initial configuration of the server.</p>	

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?	<p>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.</p> <p>Note: If capacity limitations forcing the need to use fewer disks or file systems in the environment and workloads cannot be separated, then the following can be combined on the same physical disks or file systems:</p> <ul style="list-style-type: none"> • Db2 Database active log and Database disks together • Db2 Archive log and Storage Protect storage pool disks 	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.	

Question	Tasks, characteristics, options, or settings	More information
<p>Are you setting the logs to a size that adequately supports the workload?</p>	<p>If you are not sure about the workload, use the largest size that you can.</p> <p>Active log</p> <p>Configure the server to have a maximum active log size that is appropriate for the size of your deployment. For example, for a small to medium sized deployment, you can configure a maximum active log size of 128 GB by setting the ACTIVELOGSIZE server option to a value of 131072. For more information about active log size, go to the IBM Storage Protect Blueprints and locate the Blueprint for your operating system.</p> <p>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.</p> <p>Archive log</p> <p>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.</p>	<ul style="list-style-type: none"> For information about sizing when you use data deduplication, see Checklist for data deduplication.
<p>Are you defining an archive failover log? Are you placing this log on a disk that is separate from the archive log?</p>	<p>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.</p>	<p>Use the ARCHFAILOVERLOGDIRECTORY server option to specify the location of the archive failover log.</p> <p>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.</p>
<p>If you are mirroring the active log, are you using only one type of mirroring?</p>	<p>You can mirror the log by using one of the following methods. Use only one type of mirroring for the log.</p> <ul style="list-style-type: none"> Use the MIRRORLOGDIRECTORY option that is available for the IBM Storage Protect server to specify a mirror location. Use mirroring in the disk system hardware. 	<p>If you mirror the active log, ensure that the disks for both the active log and the mirror copy have equal speed and reliability.</p> <p>For more information, see Configuring and tuning the recovery log.</p>

Planning for directory-container and cloud-container storage pools

Review how your server is set up to ensure optimal performance when using directory-container and cloud-container storage pools.

Question	Tasks, characteristics, options, or settings	More information
<p>Measured in terms of input/output operations per second (IOPS), are you using fast disk storage for the IBM Storage Protect database?</p>	<p>Use a high-performance disk for the database. Use solid-state drive technology for data deduplication processing.</p> <p>Ensure that the database has a minimum capability of 3000 IOPS. For each TB of data that is backed up daily (before data deduplication and compression), add 1000 IOPS to this minimum.</p> <p>For example:</p> <ul style="list-style-type: none"> • For daily ingest of data of 1 TB, the server needs 4000 IOPS. • For daily ingest of data of 10 TB, the server needs 13000 IOPS. • For daily ingest of data of 10 TB - 30 TB, the server needs 13000 IOPS - 33000 IOPS. • For daily ingest of data of 30 TB - 100 TB, the server needs 33000 IOPS - 100300 IOPS. <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> $3000 \text{ IOPS minimum} + 30000 (30 \text{ TB} \times 1000 \text{ IOPS}) = 33000 \text{ IOPS}$ </div>	<p>For recommendations about disk selection, see "Planning for server database disks."</p> <p>For more information about IOPS, see the IBM Storage Protect Blueprints and locate the Blueprint for your operating system.</p>

Question	Tasks, characteristics, options, or settings	More information
Do you have enough memory for the size of your database?	<p>Use a minimum of 40 GB of system memory for IBM Storage 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.</p> <p>Monitor memory usage regularly to determine whether more memory is required.</p> <p>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:</p> <ul style="list-style-type: none"> • For daily consumption of data up to 1 TB, you need 16 - 24 GB of memory. • For daily ingest of data up to 10 TB, you need 64 GB of memory. • For daily ingest of data up to 10 - 30 TB, you need 192 GB of memory. • For daily ingest of data up to 100 TB, you need 384 GB of memory. 	Memory requirements
Have you properly sized the storage capacity for the database active log and archive log?	<p>Configure the server to have a minimum active log size of 128 GB by setting the ACTIVELOGSIZE server option to a value of 131072.</p> <p>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.</p> <p>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.</p> <p>Define space for the archive failover log by using the ARCHFAILOVERLOGDIRECTORY server option.</p>	For more information about sizing for your system, see the IBM Storage Protect Blueprints and locate the Blueprint for your operating system.

Question	Tasks, characteristics, options, or settings	More information
Is compression enabled for the archive log and database backups?	<p>Enable the ARCHLOGCOMPRESS server option to save storage space.</p> <p>This compression option is different from inline compression. Inline compression is enabled by default with IBM Storage Protect 7.1.5 and later.</p> <p>Restriction: Do not use this option if the amount of backed up data exceeds 6 TB per day.</p>	For more information about compression for your system, see the IBM Storage Protect Blueprints and locate the Blueprint for your operating system.
<p>Are the IBM Storage Protect database and logs on separate disk volumes (LUNs)?</p> <p>Is the disk that is used for the database configured according to best practices for a transactional database?</p>	The database must not share disk volumes with IBM Storage 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 Storage 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.	<ul style="list-style-type: none"> • Data Deduplication FAQ • IBM Storage Protect Blueprints
Did you allocate enough storage space for the database?	<p>For a rough estimate, plan for 100 GB of database storage for every 25 TB of data that is to be protected in deduplicated storage pools. <i>Protected data</i> is the amount of data before data deduplication, including all versions of objects stored.</p> <p>For database backup operations with a large number of small files, where the average size of the file is less than 512 KB, you need more database space. For smaller object sizes, plan on 100 GB of database space for every 10 TB stored.</p> <p>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.</p>	The optimal IBM Storage Protect environment is set up by using the IBM Storage Protect Blueprints .

Question	Tasks, characteristics, options, or settings	More information
Have you estimated storage pool capacity to configure enough space for the size of your environment?	<p>You can estimate capacity requirements for a deduplicated storage pool by using the following technique:</p> <ol style="list-style-type: none">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.	<p>For an example of using this technique, see Data Deduplication FAQ.</p>

Question	Tasks, characteristics, options, or settings	More information
Have you distributed disk I/O over many disk devices and controllers?	<p>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.</p> <p>Set the <i>DB2_PARALLEL_IO</i> registry variable to enable parallel I/O for each table space used if the containers in the table space span multiple physical disks.</p> <p>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.</p> <p>Use the following guidelines to create file systems:</p> <ul style="list-style-type: none"> • For daily ingest of data less than 10 TB, you need 8 or more file systems. • For daily ingest of data of 10 TB - 30 TB, you need 12 or more file systems. • For daily ingest of data unto 100 TB, you need 32 or more file systems. 	<p>For guidelines about setting up storage pools, see "Planning for storage pools in DISK or FILE device classes."</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2® registry variables.</p>
Have you scheduled daily operations based on your backup strategy?	<p>The best practice sequence of operations is in the following order:</p> <ol style="list-style-type: none"> 1. Client backup 2. Storage pool protection 3. Data replication 4. Database backup 5. Expire inventory 	<ul style="list-style-type: none"> • Scheduling data deduplication and node replication processes • Daily operations for directory-container storage pools
Have you scheduled audit operations to identify corrupted files in storage pools?	<p>To schedule audit operations, use the DEFINE STGRULE command and specify the ACTIONTYPE=AUDIT parameter.</p> <p>As a best practice, to ensure that audit operations run continuously, do not specify the DELAY parameter.</p>	

Question	Tasks, characteristics, options, or settings	More information
Do you have enough storage to manage the IBM Db2 lock list?	<p>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.</p> <p>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.</p>	For information about tuning the Db2 LOCKLIST parameter, see .Tuning server-side data deduplication
Is sufficient bandwidth available to transfer data to an IBM Storage Protect server?	To transfer data to an IBM Storage Protect server, use client-side or server-side data deduplication and compression to reduce the bandwidth that is required.	For more information, see the enablededup client option.
Have you determined how many storage pool directories to assign to each storage pool?	<p>Assign directories to a storage pool by using the DEFINE STGPOOLDIRECTORY command.</p> <p>Create multiple storage pool directories and ensure that each directory is backed up to a separate disk volume (LUN).</p>	
Did you allocate enough disk space in the cloud-container storage pool?	<p>To prevent backup failures, ensure that the local directory has enough space. Use the following list as a guide for optimal disk space:</p> <ul style="list-style-type: none"> • 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. • For Flash or SSD based systems with fast network connections to high-performance cloud systems refer to the IBM Storage Protect Blueprints for the latest guidance on disk recommendations. 	

Question	Tasks, characteristics, options, or settings	More information
Have you bench marked the performance of your cloud container pool cloud cache?	To prevent disk hot spots, ensure that the local directory has a single storage pool directory and file system for the cloud cache.	For more information about optimizing backup operations, refer to Sizing a cloud cache to optimize backup operations.
Did you select the appropriate type of local storage?	<p>Ensure that data transfers from local storage to cloud finish before the next backup cycle starts.</p> <p>Tip: Data is removed from local storage soon after it moves to the cloud.</p> <p>Use the following guidelines:</p> <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> – 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. 	

Question	Tasks, characteristics, options, or settings	More information
<p>For cloud-container storage pools, have you specified the total maximum number of parallel processes for the storage tiering rule and each of its subrules?</p>	<p>To specify the maximum number of parallel processes, issue the DEFINE STGRULE command and specify the MAXPROCESS parameter. The default value is 8. For example, if the default value of 8 is specified, and the storage rule has four subrules, the storage rule can run eight parallel processes and each of its subrules can run eight parallel processes.</p> <p>For optimal throughput, use the following maximum number of parallel processes for small, medium, and large Blueprint systems:</p> <ul style="list-style-type: none"> • Small system: 10 processes • Medium system: 25 processes • Large system: 35-50 processes 	
<p>For cloud-container storage pools, have you defined multiple Accesser® endpoints if you are using an on-premises IBM Cloud Object Storage system with IBM Storage Protect?</p>	<p>To optimize performance, define access for the following number of Accessers for small, medium, and large blueprint systems, depending on your data ingestion requirements:</p> <ul style="list-style-type: none"> • Small system: 1 Accesser • Medium system: 2 Accessers • Large system: 3-4 Accessers 	<p>For more information, see the IBM Storage Protect Cloud Blueprints.</p>

Question	Tasks, characteristics, options, or settings	More information
<p>For cloud-container storage pools, have you defined multiple Accesser endpoints if you are using an on-premises IBM Cloud Object Storage system with IBM Storage Protect?</p>	<p>Generally, the following Ethernet capability is required to connect to private IBM Cloud Object Storage endpoints for small, medium, and large Blueprint systems:</p> <ul style="list-style-type: none"> • Small system: 1 Gbit • Medium system: 5 Gbit • Large system: 10 Gbit <p>Tip: Depending on client data ingestion and simultaneous data transfer to object storage, you might require more than one 10 Gbit Ethernet network.</p> <p>When you configure the Ethernet connection, work with a network administrator and consider the following factors:</p> <ul style="list-style-type: none"> • The Ethernet capability of the server • The nature of the network between the server and the IBM Cloud Object Storage endpoint • The final ingestion point on object storage via a cloud-container storage pool 	

Planning for daily operations for directory-container storage pools

Schedule daily operations for the server depending on the type of storage pool that you use. You can complete specific tasks with directory-container storage pools.

About this task

The following image illustrates how IBM Storage Protect tasks fit into the daily schedule.

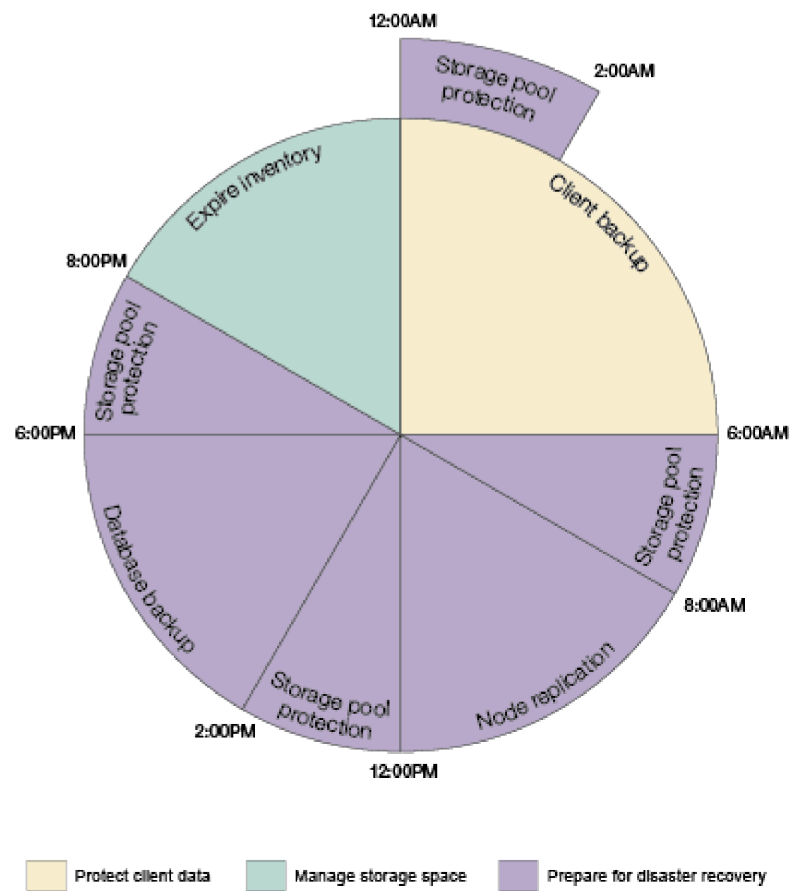


Figure 1. Daily schedule of operations for directory-container storage pools

You can schedule daily activities for IBM Storage Protect by using the Operations Center. The Operations Center creates the storage pool protection schedules when you use the wizards to configure replication or add a directory-container storage pool. You can also use the Operations Center to schedule client backups.

To manually create a schedule for daily operations, use the **DEFINE SCHEDULE** command. To plan the daily schedules for server maintenance tasks, refer to "Tuning the schedule for daily operations" in IBM Documentation.

Procedure

1. Perform an incremental backup of all clients on the network by using the **incremental** client command or use another supported method for client backup operations.
2. Create a DR copy of the IBM Storage Protect database by using the **BACKUP DB** command.
3. Protect data in directory-container storage pools to reduce node replication time by using the **PROTECT STGPPOOL** command. Protect storage pools at regular intervals during the daily schedule.
4. Perform node replication to create a secondary copy of the client data on another IBM Storage Protect server by using the **REPLICATE NODE** command.
5. Remove objects that exceed their allowed retention period by using the **EXPIRE INVENTORY** command.

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?	<p>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.</p> <p>The IBM Storage Protect server reads and writes to storage pools predominantly in 256 KB blocks.</p> <p>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.</p>	<p>For more information, see Analyzing the basic performance of disk systems.</p>

Question	Tasks, characteristics, options, or settings	More information
Did you allocate enough storage space for the database?	<p>For a rough estimate, the following database size guidelines are based on the small, medium, and large blueprint systems to allow for database growth:</p> <ul style="list-style-type: none"> • Small system: At least 2 TB • Medium system: At least 4 TB • Large system: At least 8 TB <p>Tip: You might need more memory based on the amount of data that must be protected, the number of files that are stored, and whether you use data deduplication. With data deduplication, the load on the database becomes greater because there are frequent queries to the database to determine what deduplicated extents are on the server.</p> <p>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.</p> <p>If you have several hundred TB of protected data, or if you are backing up multiple TBs of data daily, the starting size for the database must be at least 1 TB. Use the IBM Storage Protect to size the database for your system.</p>	<p>The optimal IBM Storage Protect environment is set up by using the IBM Storage Protect Blueprints.</p> <p>For information about the minimum amount of memory you must allocate on the server to complete operations, based on the database size, see Memory requirements.</p>
Is the disk configured to use read and write cache?	Use more cache for better performance.	
Do you need to backup the IBM Storage Protect database to cloud object storage?	<p>You can back up a database to, and restore a database from, cloud object storage for disaster recovery purposes.</p> <p>You can tune object storage endpoints, IBM Cloud Object Storage Accessers, network bandwidth, and data streams to ensure that database backup operations run efficiently.</p>	<p>Tuning database backups to cloud object storage.</p>

Question	Tasks, characteristics, options, or settings	More information
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 system 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.
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: <ul style="list-style-type: none">• The MAXSESSIONS server option, which is the maximum number of IBM Storage Protect sessions that can run concurrently.• The MAXNUMP 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 MAXNUMP =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 MAXNUMP parameter for client nodes.

Question	Tasks, characteristics, options, or settings	More information
For storage pools that use DISK device classes, have you determined how many storage pool volumes to put on each file system?	<p>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.</p> <p>If you are not using RAID, then configure one file system per physical disk, and define one storage pool volume for each file system.</p> <p>If you are using RAID 5 with $n + 1$ volumes, configure the storage in one of the following ways:</p> <ul style="list-style-type: none"> • 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?	<p>Ensure that each file system is on a different LUN on the disk system.</p> <p>Typically, having 10 - 30 file systems is a good goal, but ensure that the file systems are no smaller than approximately 250 GB.</p>	<p>For details, see the following topics:</p> <ul style="list-style-type: none"> • Tuning disk storage for the server • Tuning and configuring storage pools and volumes
Have you scheduled audit operations to identify corrupted files in storage pools?	To schedule audit operations, use the DEFINE SCHEDULE command to run AUDIT VOLUME FIX=NO commands.	

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 Storage 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 5. Storage technology types for IBM Storage 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: <ul style="list-style-type: none"> – You are using IBM Storage Protect data deduplication. – You are backing up more than 8 TB of new data daily. 	If you place the IBM Storage 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: <ul style="list-style-type: none"> – 15k rpm disk – Fibre Channel or serial-attached SCSI (SAS) interface 	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> – The server does not do data deduplication. – The server does not do data 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: <ul style="list-style-type: none"> – The server does not do data deduplication. – The server does not do data 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: <ul style="list-style-type: none"> – 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: <ul style="list-style-type: none"> – 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: <ul style="list-style-type: none"> – 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 5. Storage technology types for IBM Storage 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: <ul style="list-style-type: none"> – 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 Storage Protect solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of IBM Storage 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 data replication. Avoid Serial Advanced Technology Attachment (SATA) and Parallel Advanced Technology Attachment (PATA) disks. For details and more tips, see the following topics: <ul style="list-style-type: none"> – "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 84861 . Heavier workloads require more than the minimum requirements. Advanced features such as data deduplication and data 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.	<p>Keep all IBM Storage 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:</p> <ul style="list-style-type: none"> – "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.	<p>Place each directory on a LUN that is isolated from other LUNs and from other applications.</p> <p>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.</p> <p>See "Planning for server database disks."</p>
If you are using data deduplication, data replication, or both, follow the guidelines for database configuration and other items.	<p>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:</p> <ul style="list-style-type: none"> – 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.	<p>Review the information in Optimal number and size of volumes for storage pools that use disk to help you to determine volume size.</p> <p>Configure storage pool devices and file systems based on throughput requirements, not only on capacity requirements.</p> <p>Isolate the storage devices that are used by IBM Storage Protect from other applications that have high I/O, and ensure that there is enough throughput to that storage.</p> <p>For more details, see Checklist for storage pools on DISK or FILE.</p>
Schedule IBM Storage Protect client operations and server maintenance activities to avoid or minimize overlap of operations.	<p>For more details, see the following topics:</p> <ul style="list-style-type: none"> – Tuning the schedule for daily operations – Checklist for server configuration
Monitor operations constantly.	<p>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.</p>

Minimum system requirements for the IBM Storage Protect server

Before you install an IBM Storage Protect server on a Windows operating system, review the hardware and software requirements.

Hardware and software requirements for the IBM Storage Protect server installation

The optimal IBM Storage Protect environment is set up with data deduplication by using the [IBM Storage Protect Blueprints](#).

For the most current information about IBM Storage Protect system requirements, see [technote 1243309](#).

IBM Installation Manager

IBM Storage 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 Storage Protect. It must remain installed on the system so that IBM Storage 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 Storage Protect offering contains all of the media that IBM Installation Manager requires to install IBM Storage Protect.

Package

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

The IBM Storage Protect package contains the following components:

- IBM Installation Manager installation program
- IBM Storage Protect offering

Package group

A set of packages that share a common parent directory.

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

Repository

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

The IBM Storage 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 Storage 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 Storage Protect server. You can also use them to keep track of names and user IDs.

Restriction: If you are using a File Allocation Table (FAT or FAT32) or a New Technology File System (NTFS) format, you cannot specify the root directory of that system as the location of a database directory or log directory. Instead, you must create one or more subdirectories within the root directory. Then, create the database directories and log directories within the subdirectories.

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 <i>instance user ID</i> for the server, which is the ID you use to start and run the IBM Storage Protect server		
The <i>home directory</i> for the server, which is the directory that contains the instance user ID		
The database instance name		
The <i>instance directory</i> 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 Storage Protect includes managing resources such as the database, the recovery log and the shared resource area.

Before you begin

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 Storage Protect database is 8 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:

$$100,000 \text{ files} \times 200 \text{ bytes} = 19 \text{ MB}$$

- b) Calculate the number of copy storage-pool files.

All primary storage pools are backed up to the copy storage pool:

$$(1,500,000 + 100,000 + 200,000) \times 200 \text{ bytes} = 343 \text{ MB}$$

- 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:

$$500,000 \times 200 \text{ bytes} = 95 \text{ MB}$$

- d) Calculate the number of deduplicated files.

Assume that a deduplicated storage pool contains 50,000 files:

$$50,000 \times 200 \text{ bytes} = 10 \text{ MB}$$

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.

$$(1.8 + 0.5) \times 50\% = 1.2 \text{ GB}$$

4. Calculate the total amount of database space that is required for the client. The total is approximately 3.5 GB:

$$1.8 + 0.5 + 1.2 = 3.5 \text{ GB}$$

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:

$$500 \times 3.5 = 1.7 \text{ TB}$$

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 Storage 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 Storage 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 Storage Protect server must run with limited memory. The database uses the IBM Storage Protect server main memory to run database operations. However, if there is insufficient memory available, the IBM Storage 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 Storage 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 Storage 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 Storage 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 Storage 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 Storage 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.

<i>Table 6. Basic client-store operations</i>		
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	<p>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.</p> <p>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.</p>
Active log: Suggested size	19.5 GB ¹	<p>Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.</p> <p>$(300 \text{ clients} \times 4096 \text{ files stored during each transaction} \times 3053 \text{ bytes for each file}) \div 1,073,741,824 \text{ bytes} = 3.5 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>$3.5 + 16 = 19.5 \text{ GB}$</p>
Archive log: Suggested size	58.5 GB ¹	<p>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.</p> <p>$3.5 \times 3 = 10.5 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 48 GB:</p> <p>$10.5 + 48 = 58.5 \text{ GB}$</p>
<p>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.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

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:

$$\begin{array}{l} \text{number of clients} \times \text{sessions for each client} \times \text{files stored} \\ \text{during each transaction} \times \text{log space needed for each file} \end{array}$$

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

<i>Table 7. Multiple client sessions</i>			
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	<p>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.</p> <p>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.</p>

Table 7. Multiple client sessions (continued)

Item	Example values		Description
Active log: Suggested size	26.5 GB ¹	51 GB ¹	<p>The following calculation was used for 300 clients. One GB equals 1,073,741,824 bytes.</p> $(300 \text{ clients} \times 3 \text{ sessions for each client} \times 4096 \text{ files stored during each transaction} \times 3053 \text{ bytes for each file}) \div 1,073,741,824 = 10.5 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $10.5 + 16 = 26.5 \text{ GB}$ <p>The following calculation was used for 1000 clients. One GB equals 1,073,741,824 bytes.</p> $(1000 \text{ clients} \times 3 \text{ sessions for each client} \times 4096 \text{ files stored during each transaction} \times 3053 \text{ bytes for each file}) \div 1,073,741,824 = 35 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $35 + 16 = 51 \text{ GB}$
Archive log: Suggested size	79.5 GB ¹	153 GB ¹	<p>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:</p> $10.5 \times 3 = 31.5 \text{ GB}$ $35 \times 3 = 105 \text{ GB}$ <p>Increase those amounts by the suggested starting size of 48 GB:</p> $31.5 + 48 = 79.5 \text{ GB}$ $105 + 48 = 153 \text{ GB}$

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.

<i>Table 8. Simultaneous write operations</i>		
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	<p>3053 bytes plus 200 bytes for each copy storage pool.</p> <p>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.</p> <p>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.</p>
Active log: Suggested size	20 GB ¹	<p>Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.</p> <p>$(300 \text{ clients} \times 4096 \text{ files stored during each transaction} \times 3453 \text{ bytes for each file}) \div 1,073,741,824 \text{ bytes} = 4.0 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>$4 + 16 = 20 \text{ GB}$</p>
Archive log: Suggested size	60 GB ¹	<p>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:</p> <p>$4 \text{ GB} \times 3 = 12 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 48 GB:</p> <p>$12 + 48 = 60 \text{ GB}$</p>
<p>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.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

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 Storage 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.

<i>Table 9. Full database backups</i>		
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	<p>3053 bytes for each file plus 200 bytes for each copy storage pool.</p> <p>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.</p> <p>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.</p>

Table 9. Full database backups (continued)		
Item	Example values	Description
Active log: Suggested size	20 GB ¹	<p>Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.</p> $(300 \text{ clients} \times 4096 \text{ files per transaction} \times 3453 \text{ bytes per file}) \div 1,073,741,824 \text{ bytes} = 4.0 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $4 + 16 = 20 \text{ GB}$
Archive log: Suggested size with a full database backup every day	60 GB ¹	<p>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:</p> $4 \text{ GB} \times 3 = 12 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $12 + 48 = 60 \text{ GB}$
Archive log: Suggested size with a full database every week	132 GB ¹	<p>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:</p> $(4 \text{ GB} \times 3) \times 7 = 84 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $84 + 48 = 132 \text{ GB}$
<p>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.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

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 10. 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: $(800 \text{ GB} \div 700 \text{ KB}) = 1,198,372 \text{ bits}$ The following calculation was used for a 4 TB object: $(4 \text{ TB} \div 700 \text{ KB}) = 6,135,667 \text{ bits}$
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.
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: $(23.3 \text{ GB} + 1.7 \text{ GB}) \times 2 = 50 \text{ GB}$ Increase that amount by the suggested starting size of 16 GB: $50 + 16 = 66 \text{ GB}$ The following calculation was used for multiple transactions and a 4 TB object: $(23.3 \text{ GB} + 8.6 \text{ GB}) \times 2 = 63.8 \text{ GB}$ Increase that amount by the suggested starting size of 16 GB: $63.8 + 16 = 79.8 \text{ GB}$

Table 10. Average duplicate-extent size of 700 KB (continued)

Item	Example values		Description
Archive log: Suggested size	198 GB ¹	239.4 GB ¹	<p>Multiply the estimated size of the active log by a factor of 3.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $50 \text{ GB} \times 3 = 150 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $150 + 48 = 198 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $63.8 \text{ GB} \times 3 = 191.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $191.4 + 48 = 239.4 \text{ GB}$

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 11. 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.
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	<p>Using the average extent size, these calculations represent the total number of extents for a given object.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(800 \text{ GB} \div 256 \text{ KB}) = 3,276,800 \text{ bits}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(4 \text{ TB} \div 256 \text{ KB}) = 16,777,216 \text{ bits}$

Table 11. Average duplicate-extent size of 256 KB (continued)			
Item	Example values		Description
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	<p>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.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 4.5 \text{ GB}) \times 2 = 55.6 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $55.6 + 16 = 71.6 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 23.4 \text{ GB}) \times 2 = 93.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $93.4 + 16 = 109.4 \text{ GB}$
Archive log: Suggested size	214.8 GB ¹	328.2 GB ¹	<p>The estimated size of the active log multiplied by a factor of 3.</p> <p>The following calculation was used for an 800 GB object:</p> $55.6 \text{ GB} \times 3 = 166.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $166.8 + 48 = 214.8 \text{ GB}$ <p>The following calculation was used for a 4 TB object:</p> $93.4 \text{ GB} \times 3 = 280.2 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $280.2 + 48 = 328.2 \text{ GB}$

Table 11. Average duplicate-extent size of 256 KB (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 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 Storage 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 Storage 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 Storage 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 DSMSEV 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 Storage Protect server is created with the DSMSEV 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.
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:

`C:\Program Files\IBM\Installation Manager\eclipse\tools`

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

`imcl.exe -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 Storage 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.

Database instance name on Windows

The database instance name is the name of the server instance as it appears in the registry.

For example: `Server1`

Instance directory on Windows

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 use a name that includes the name of the server instance as it appears (or will appear) in the registry. Default server instance names have the form `Serverx`.

For example: `C:\tsm\server1`

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 Storage Protect, and is used for operations that involve communication among multiple IBM Storage 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 Storage 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,

```
TUCSON_SERVER1
TUCSON_SERVER2
```

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:

```
f:\server1\archlog
```


Chapter 2. Installing the server components

To install the IBM Storage Protect server components, you can use either the installation wizard or the command line in console mode.

About this task

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

- server

Tip: The database (IBM 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 Storage Protect for SAN
- Operations Center
- Open Snap Store Manager (OSSM)

Restriction: The system where you plan to install OSSM must be running on a Linux® x86_64 operating system.

Allow approximately 15 - 30 minutes to install a server, using this guide.

Obtaining the installation package

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

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 Storage Protect [technote 588021](#)
 - IBM Storage Protect Extended Edition [technote 588023](#)
 - IBM Storage Protect for Data Retention [technote 588025](#)
 - b. Change to the directory where you placed the executable file.

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

```
package_name.exe
```

where *package_name* is like this example: *8.1.x.000-IBM-SPSRV-WindowsX64.exe*

3. Select one of the following methods of installing IBM Storage Protect:
 - [“Installing IBM Storage Protect by using the installation wizard” on page 64](#)
 - [“Installing IBM Storage Protect by using console mode” on page 65](#)
 - [“Installing IBM Storage Protect in silent mode” on page 67](#)
4. After you install IBM Storage 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 Storage 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.
- Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

Install IBM Storage Protect by using this method:

Option	Description
Installing the software from a downloaded package:	<ol style="list-style-type: none">a. Change to the directory where you downloaded the package.b. Start the installation wizard by issuing the following command:<div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"><code>install.bat</code></div> <p>Or, in the directory where the installation files were extracted, double-click the <code>install.bat</code> file.</p> <p>Tip: To upgrade the instance, you must select Y to Do you update the instance? question. If you select No, the instance is ignored and deleted. You must then recreate and upgrade the instance manually after the upgrade is completed.</p>

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 [Chapter 3, “Taking the first steps after you install IBM Storage Protect,” on page 71](#) to learn about configuring your server.
- If a native device driver is available on Windows for the tape drives or medium changers that you plan to use, use the native device driver. If a native device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Storage Protect device driver by issuing the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is `C:\Program Files\Tivoli\TSM\device\drivers`.

Installing IBM Storage Protect by using console mode

You can install IBM Storage 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.
- Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

Install IBM Storage Protect by using this method:

Option	Description
Installing the software from a downloaded package:	<p>a. Change to the directory where you downloaded the package.</p> <p>b. Start the installation wizard in console mode by issuing the following command:</p> <pre>install.bat -c</pre> <p>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.</p>
Upgrading the software by using the console mode	<p>a. Use cd command to change the current directory to the directory where you extracted the software package. For example:</p> <pre>cd /code/software/server</pre> <p>b. Start the installation wizard in console mode by issuing the following command:</p> <pre>./install.sh -c</pre> <p>Output:</p> <pre>Preprocessing the input. =====> IBM Installation Manager Select: 1. Install - Install software packages 2. Update - Find and install updates and fixes to installed software packages 3. Modify - Change installed software packages 4. Roll Back - Revert to an earlier version of installed software packages 5. Uninstall - Remove installed software packages Other Options: L. View Logs S. View Installation History V. View Installed Packages ----- P. Preferences ----- A. About IBM Installation Manager ----- X. Exit Installation Manager -----></pre> <p>c. Select 2. Update - Find and install updates and fixes to installed software packages</p>

Option	Description
	<p>Output:</p> <pre> =====> IBM Installation Manager> Update Select a package group to update: 1. [X] IBM Storage Protect Details of package group IBM Storage Protect: Package Group Name : IBM Storage Protect Shared Resources Directory : /opt/IBM/IBMIMShared Installation Directory : /opt/tivoli/tsm Translations : English Architecture : 64-bit Other Options: U. Update All A. Unselect All N. Next, C. Cancel -----> [N]</pre> <p>d. To upgrade, ensure that IBM Storage Protect is selected. Select N to continue.</p> <p>Output:</p> <pre> =====> IBM Installation Manager> Update> Packages Package group: IBM Storage Protect Update packages: 1. [X] IBM Storage Protect server 8.1.13.20211106_0212 2. [X] IBM Storage Protect license 8.1.13.20211106_0206 3. [X] IBM Storage Protect Operations Center 8.1.13000.20211105_1031 4. [X] Open Snap Store Manager 8.1.13.20211106_0205 Other Options: B. Back, U. Update C. Cancel</pre> <p>e. Ensure that you select all the products that require to be upgraded. Select U to start the upgrade.</p> <p>Tip: To upgrade the instance, you must select Y to Do you update the instance? question. If you select No, the instance is ignored and deleted. You must then recreate and upgrade the instance manually after the upgrade is completed.</p>

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:
C:\ProgramData\IBM\Installation Manager\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 [Chapter 3, “Taking the first steps after you install IBM Storage Protect,” on page 71](#) to learn about configuring your server.
- If a native device driver is available on Windows for the tape drives or medium changers that you plan to use, use the native device driver. If a native device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Storage Protect device driver by issuing the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is C:\Program Files\Tivoli\TSM\device\drivers.

Installing IBM Storage 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 Storage Protect components.

update_response_sample.xml

Use this file to upgrade the IBM Storage 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.bat -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:

```
C:\ProgramData\IBM\Installation Manager\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 [Chapter 3, “Taking the first steps after you install IBM Storage Protect,”](#) on page 71 to learn about configuring your server.
- If a native device driver is available on Windows for the tape drives or medium changers that you plan to use, use the native device driver. If a native device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Storage Protect device driver by issuing the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is `C:\Program Files\Tivoli\TSM\device\drivers`.

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 server messages and help: LANGUAGE AMENG.

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 12. Server languages for Windows	
Language	LANGUAGE option value
Chinese, Simplified	chs
Chinese, Traditional	cht
English	ameng
French	fra
German	deu
Italian	ita
Japanese (Shift-JIS)	jpn
Korean	kor
Portuguese, Brazilian	ptb
Russian	rus
Spanish	esp

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 Storage Protect.

About this task

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 `ita` locale, set the LANGUAGE option to `ita`. See [“Server language locales” on page 68](#).

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 Storage 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.

Installing Open Snap Store Manager

On the IBM Storage Protect server, use the IBM Installation Manager graphical wizard to install the Open Snap Store Manager (OSSM) server component. You can install OSSM using the installation wizard or console mode.

Before you begin

Ensure that the following prerequisites are met:

- The system where you plan to install OSSM must be running on a Linux® x86_64 operating system.
- You must have system privileges on the IBM Storage Protect server.
- The IBM Storage Protect server and the OSSM component must be installed on the same system:
 - If you are upgrading the IBM Storage Protect server: Upgrade the server and other non-OSSM components first. Then, run the installation process again and install only OSSM.
 - If you are installing the IBM Storage Protect server for the first time: Select all components that you plan to install, including the server and OSSM, at the same time.

Procedure

To install the OSSM component, complete the following steps:

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

```
./install.sh
```

This opens the IBM Installation Manager screen.

3. Select the OSSM component in the IBM Installation Manager and follow the steps in the wizard to complete the installation.

What to do next

For more information on configuring the OSSM instance and backing up VMware data to the OSSM server, refer to [Backing up VMware data to the Open Snap Store Manager](#).

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

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

About this task

1. Create the directories and user ID for the server instance. See [“Creating the user ID and directories for the server instance”](#) on page 71.
2. Configure a server instance. Select one of the following options:
 - Use the configuration wizard, the preferred method. See [“Configuring IBM Storage Protect by using the configuration wizard”](#) on page 73.
 - Manually configure the new instance. See [“Configuring the server instance manually”](#) on page 75. Complete the following steps during a manual configuration.
 - a. Set up your directories and create the IBM Storage Protect instance. See [“Creating the server instance”](#) on page 75.
 - 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 76.
 - c. Issue the **DSMSERV FORMAT** command to format the database. See [“Formatting the database and log”](#) on page 78.
3. Configure options to control when database reorganization runs. See [“Configuring server options for server database maintenance”](#) on page 79.
4. Start the server instance if it is not already started.

See [“Starting the server instance on Windows systems”](#) on page 80.
5. Register your license. See [“Registering licenses”](#) on page 84.
6. Prepare your system for database backups. See [“Preparing the server for database backup operations”](#) on page 85.

Tip: Back up files that are not saved during a database backup operation. For example, the following files and directories are required to restore a server:

 - Server options file (dsmserv.opt)
 - Device configuration file (for example, devconf.dat)
 - Volume history file (for example, volhist.dat)
 - Master encryption key files (dsmkeydb.kdb or dsmkeydb.sth)
 - Server certificate and private key files (cert.kbd or cert.sth)
7. To facilitate troubleshooting in case of any future issues, ensure that sufficient space is allocated for a core dump. For more information, see [technote 6357399](#).
8. Monitor the server. See [“Monitoring the server”](#) on page 86.

Creating the user ID and directories for the server instance

Create the user ID for the IBM Storage 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 42.

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 that will be the owner of the IBM Storage Protect server instance. A user ID can own more than one IBM Storage Protect server instance. Identify the user account that will own the server instance.

When the server is started as a Windows service, this account is the one that the service will log on to. The user account must have administrative authority on the system. One user account can own more than one server instance.

If you have multiple servers on one system and want to run each server with a different user account, create a new user account in this step.

Create the user ID.

Restriction: The user ID must comply with the following rule:

In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID must be 30 characters or less, and 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.

- a. Use the following operating system command to create the user ID:

```
net user user_ID * /add
```

You are prompted to create and verify a password for the new user ID.

- b. Issue the following operating system commands to add the new user ID to the Administrators groups:

```
net localgroup Administrators user_ID /add
net localgroup DB2ADMNS user_ID /add
```

2. Create directories that the server requires.

Create empty directories for each item in the table and ensure that the new user ID you just created has read/write permission to the directories. The database, archive log, and active log must reside on different physical volumes.		
Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> 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)	<code>mkdir d:\tsm\server1</code>	
The database directories	<code>mkdir d:\tsm\db001</code> <code>mkdir e:\tsm\db002</code> <code>mkdir f:\tsm\db003</code> <code>mkdir g:\tsm\db004</code>	
Active log directory	<code>mkdir h:\tsm\log</code>	
Archive log directory	<code>mkdir i:\tsm\archlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir j:\tsm\logmirror</code>	

Create empty directories for each item in the table and ensure that the new user ID you just created has read/write permission to the directories. The database, archive log, and active log must reside on different physical volumes. *(continued)*

Item	Example commands for creating the directories	Your directories
Optional: Secondary archive log directory (failover location for archive log)	mkdir k:\tsm\archlogfailover	

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 Storage Protect server

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

About this task

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

- Use the IBM Storage Protect configuration wizard on your local system. See [“Configuring IBM Storage Protect by using the configuration wizard”](#) on page 73.
- Manually configure the new IBM Storage Protect instance. See [“Configuring the server instance manually”](#) on page 75. Complete the following steps during a manual configuration.
 1. Set up the directories and create the IBM Storage Protect instance. See [“Creating the server instance”](#) on page 75.
 2. Create a new server options file by copying the sample file in order to set up communications between the IBM Storage Protect server and clients. See [“Configuring server and client communications”](#) on page 76.
 3. Issue the DSMSERV FORMAT command to format the database. See [“Formatting the database and log”](#) on page 78.

Configuring IBM Storage 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 Storage Protect server program.

Before you begin

Before you use the configuration wizard, you must complete all preceding steps to prepare for the configuration. These steps include installing IBM Storage 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:
 - Start the Remote Registry service:
 - a. Click **Start > Administrative Tools > Services**.
 - b. In the **Services** window, select the **Remote Registry** service if it is not started, and click **Start**.

- Ensure that the following ports are not blocked by a firewall: 137, 139 and 445. Complete the following steps:
 - a. Click **Start > Control Panel > Windows Firewall**.
 - b. Select **Advanced Settings**.
 - c. Select **Inbound Rules** in the navigation pane.
 - d. Select **New Rule**.
 - e. Create a port rule for TCP ports 137, 139 and 445 to allow connections for domain and private networks.
- Configure **User Account Control**:

Access all three of the user account control configuration settings by first accessing **Local Security Policy Security** options. For more information, see [“Configuring Remote Execution Protocol on Windows” on page 74](#).

2. Start the local version of the wizard:

Either click **Start > All Programs > IBM Storage Protect > Configuration Wizard**. Or, double-click the `dsmicfgx.exe` program in *installation_directory\server*. The default directory is `C:\Program Files\Tivoli\TSM`.

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 Remote Execution Protocol on Windows

Configure remote access settings by using these procedures.

Before you begin

You must configure the User Account Control feature before you run the wizard.

Procedure

If the system is running on Windows, complete the following steps to disable User Account Control:

- a) Enable the built-in Administrator account. Complete the following steps:
 - i) Click **Control Panel > Administrative Tools > Local Security Policy**.
 - ii) Under **Security Settings**, double-click **Local Policies**.
 - iii) Double-click **Security Options**.
 - iv) Double-click the **Accounts: Administrator account status** section.
 - v) Select **Enable** and click **OK**.
- b) Disable User Account Control for all Windows administrators. Complete the following steps:
 - i) Click **Control Panel > Administrative Tools > Local Security Policy**.
 - ii) Under **Security Settings**, double-click **Local Policies**.
 - iii) Double-click **Security Options**.
 - iv) Double-click the **User Account Control: Run all administrators in Admin Approval Mode** section.
 - v) Select **Disable** and click **OK**.
- c) Disable User Account Control for the built-in Administrator account. Complete the following steps:
 - i) Click **Control Panel > Administrative Tools > Local Security Policy**.
 - ii) Under **Security Settings**, double-click **Local Policies**.
 - iii) Double-click **Security Options**.
 - iv) Double-click the **User Account Control: Admin Approval Mode for the Built-in Administrator Account** section.

v) Select **Disable** and click **OK**.

Configuring the server instance manually

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

Creating the server instance

Create an IBM Storage 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:

- Ensure that the user and the instance directory of the user exists. If there is no instance directory, you must create it.

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)
- Save a backup copy of the following files to a safe and secure location:
 - Master encryption key files (`dsmkeydb.*`)
 - Server certificate and private key files (`cert.*`)
1. Log in as an administrator and create an IBM Storage Protect instance, by using the **db2icrt** command. Enter the following command on one line. The user account that you specify becomes the user ID that owns the server (the instance user ID).

```
db2icrt -u user_account instance_name
```

For example, if the user account is *tsminst1* and the server instance is *Server1*, enter the following command:

```
db2icrt -u tsminst1 server1
```

You are prompted for the password for user ID *tsminst1*. Later, when you create and format the database, you use the instance name that you specified with this command, with the `-k` option.

2. Change the default path for the database to be the drive where the instance directory for the server is located. Complete the following steps:
 - a. Click **Start > Programs > IBM DB2 > DB2TSM1 > Command Line Tools > Command Line Processor**.
 - b. Enter `quit` to exit the command line processor.
A window with a command prompt should now be open, with the environment properly set up to successfully issue the commands in the next steps.
 - c. From the command prompt in that window, issue the following command to set the environment variable for the server instance that you are working with:

```
set db2instance=instance_name
```

The *instance_name* is the same as the instance name that you specified when you issued the **db2icrt** command. For example, to set the environment variable for the *Server1* server instance, issue the following command:

```
set db2instance=server1
```

d. Issue the command to set the default drive:

```
db2 update dbm cfg using dftdbpath instance_location
```

For example, the instance directory is *d:\tsm\server1* and the instance location is drive *d:*. Enter the command:

```
db2 update dbm cfg using dftdbpath d:
```

3. Create a new server options file.

Configuring server and client communications

After installing the server, you can set up client and server communications by specifying options in the server and client options files.

About this task

Set these server options before you start the server. When you start the server, the new options go into effect. If you modify any server options after starting the server, you must stop and restart the server to activate the updated options.

Review the server options file (*dsmserv.opt.smp*) that is located in the server instance directory to view and specify server communications options. By default, the server uses the TCP/IP and Named Pipes communication methods.

Tip: If you start the server console and see warning messages that a protocol could not be used by the server, either the protocol is not installed or the settings do not match the Windows protocol settings.

For a client to use a protocol that is enabled on the server, the client options file must contain corresponding values for communication options. In the server options file, you can view the values for each protocol.

You can specify one or more of the following communication methods:

- TCP/IP Version 4 or Version 6
- Named Pipes
- Shared memory
- Secure Sockets Layer (SSL)

Tip: You can authenticate passwords with the LDAP directory server, or authenticate passwords with the 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 Storage 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.

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.

SSLTCPSPORT

(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 Named Pipes options

The Named Pipes communication method is ideal when running the server and client on the same Windows machine. Named Pipes require no special configuration.

About this task

Here is an example of a Named Pipes setting:

```

commmethod      namedpipe
namedpipename   \\.\pipe\adsmpipe

```

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

```

commmethod tcpip
commmethod namedpipe

```

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 Storage Protect server to manage the increased requirements.

Formatting the database and log

If you configure the server manually, you must format the server database and recovery log. The database is used to store information about client data and server operations and the recovery log can be used to recover from system and media failures. Use the **DSMSERV FORMAT** utility to format and initialize the server database and recovery log. 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. Do not place the directories on file systems that might run out of space. If certain directories, such as the archive log, are no longer available or full, the server stops. See [Capacity planning](#) for more details.

Important: The installation program creates a set of registry keys. One of these keys points to the directory where a default server, named SERVER1, is created. To install another server, create a directory and use the **DSMSERV FORMAT** utility, with the **-k** parameter, from that directory. The new directory becomes the location of the server. The registry tracks the installed servers.

Setting the exit list handler

Set the **DB2NOEXITLIST** registry variable to ON for each server instance. Log on to the system by using the instance user ID and run the following command:

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

For example:

```
db2set -i server1 DB2NOEXITLIST=ON
```

Initializing the server database and recovery log

Use the **DSMSERV FORMAT** utility to format and initialize the server database, which is an IBM Db2 database, and the recovery log. For example, if the server instance directory is */tsminst1*, run the following commands:

```
cd \tsminst1
dsmserv -k server2 format dbdir=d:\tsm\db001 activelogsiz=32768
activelogdirectory=e:\tsm\activelog archlogdirectory=f:\tsm\archlog
archfailoverlogdirectory=g:\tsm\archfaillog mirrorlogdirectory=h:\tsm\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.

Creating an administrative user

After the formatting of the database and recovery log is completed, you must create an administrative user who can log in to the server and also enable the IBM Storage Protect Operations Center to connect to the server. You use the following commands in a macro to set up an administrative user:

REGISTER ADMIN

The **REGISTER ADMIN** command takes the following parameters:

```
register admin administrator_user_id administrator_user_password
```

The password must meet specific length rules. For more information, see [REGISTER ADMIN \(Register an administrator ID\)](#)

GRANT AUTH

The **GRANT AUTH** command takes the following parameters:

```
grant auth administrator_user_id classes=administrator_user_class
```


For more information, see [GRANT AUTHORITY \(Add administrator authority\)](#).

Complete the following steps to set up an administrative user:

1. Create a macro, for example, `setup.mac`.
2. Edit the macro to register an administrative user and grant system authority to the user, with the following credentials:
 - Administrative user ID: `adminadmin`
 - Password for the administrative user: `adminadmin1`

```
register admin adminadmin adminadmin1
grant auth adminadmin classes=system
```

You must create the administrative user with the **classes=system** option so that the administrative user can create other potential administrative users, for example, with limited privileges. Any of these administrative users can then connect to the IBM Storage Protect Operations Center.

3. To create the administrative user and grant system authority to this user, run the **DSMSERV** command with the **runfile** option and the macro file, for example:

```
dsmserv runfile setup.mac
```

The administrative user can then start the server instance and connect to the server to complete other required steps, such as setting up the database backup.

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 by using a text editor. 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 `c:\Program Files\Tivoli\TSM` 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.

- a) 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:

```
reorgbegin time 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.

Starting the server instance on Windows systems

In a production environment, the preferred method for starting the server is as a Windows service. In an environment where you are reconfiguring, testing, or completing maintenance tasks, start the server in the foreground or use maintenance mode.

Before you begin

Select one of the following methods for starting the server:

As a Windows service

This method is useful in a production environment. When you configure the server to run as a service, you can specify that the server starts automatically whenever the system is started.

In the foreground

This method is useful when you are configuring or testing the server. When you start the server in the foreground, IBM Storage Protect provides a special administrator user ID that is named **SERVER_CONSOLE**. All server messages are displayed in the foreground. The messages can be useful if you must debug startup problems.

In maintenance mode

This method is useful when you are completing maintenance or reconfiguration tasks. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

Procedure

Follow the instructions for your selected option:

Option	Description
Starting the server as a Windows service	<p>To start the server as a Windows service, take one of the following actions:</p> <ul style="list-style-type: none"> • If you configured the server by using the configuration wizard, complete the following steps: <ol style="list-style-type: none"> a. Configure the server to start as a Windows service by following the instructions in “Configuring the server to start as a Windows service” on page 81.

Option	Description
	<p>b. Start the server by following the instructions in “Starting the server as a Windows service” on page 82.</p> <ul style="list-style-type: none"> If you did not use the configuration wizard, create and configure the Windows service by following the instructions in “Manually creating and configuring a Windows service” on page 82.
Starting the server in the foreground	To start the server in the foreground, follow the instructions in “Starting the server in the foreground” on page 83.
Starting the server in maintenance mode	To start the server in maintenance mode, follow the instructions in “Starting the server in maintenance mode” on page 83.

Configuring the server to start as a Windows service

Before you can start the server as a Windows service, you must ensure that options and access rights are set correctly.

Before you begin

A Windows service must be created. If you configured the server by using the configuration wizard, a Windows service was created automatically. In that case, use this procedure to configure the server to start as a Windows service.

If you did not use a wizard, you must create and configure the Windows service manually by following the steps in [“Manually creating and configuring a Windows service”](#) on page 82.

Procedure

- From the Windows **Start** menu, click **Run**, type `services.msc`, and click **OK**.
- In the **Services** window, select the server instance that you want to start as a service, and click **Properties**.
For example, select **TSM INST1**, and click **Properties**.
- To ensure that the server service starts automatically, click the **General** tab. From the **Startup type** list, select **Automatic**.
- To set the user for starting the server service, click the **Log On** tab, and take one of the following actions:
 - If you plan to run the server service under the Local System account, select **Local System account** and click **OK**.
 - If you plan to run the server service under the instance user ID, take the following actions:
 - Select **This account**, and browse for the user ID that owns the server Db2 instance and has permissions for starting the server.
 - In the **Select User** window, in the **Enter the object name to select** field, enter the user ID.
 - Click **Check Names**.
 - Click **OK** twice.
- If you configured the server service to run under the Local System account, grant database access to the Local System account:
 - Log on with the user ID that was used to create the server database. This user ID is the user ID that was used to run the **DSMSERV FORMAT** utility to initialize the server database. Alternatively, if you configured the server with the **dsmicfgx** configuration wizard, this user ID is the user ID that was used to create the instance.
 - Open a Db2 command window. If the server is installed on Windows Server 2012, open the **Start** window, and click **DB2 Command Window - Administrator**.

c) In the Db2 command window, enter the following commands:

```
set DB2INSTANCE=server1
db2 connect to TSMDB1
db2 grant dbadm with dataaccess with accessctrl on database to user system
db2 grant secadm on database to user system
```

Tip: When the server service is configured to run under the Local System account, the database can be accessed by any administrator on the system. In addition, any administrator who can log on to the system can run the server.

What to do next

To start the service, follow the instructions in [“Starting the server as a Windows service”](#) on page 82.

Starting the server as a Windows service

If you are running IBM Storage Protect on a Windows operating system, you can start the server as a service.

Before you begin

A Windows service must be created. The service was created automatically if you configured the server by using the configuration wizard. If the service was created automatically, you must configure the server to start as a service by following the steps in [“Configuring the server to start as a Windows service”](#) on page 81. Then, use this procedure to start the server as a service.

If you did not use the configuration wizard to create the service, you must create and configure the service manually. Follow the steps in [“Manually creating and configuring a Windows service”](#) on page 82.

Procedure

To start the server as a Windows service, complete the following steps:

1. Log on to the server with a user ID that is in the Administrators group.
2. From the Windows **Start** menu, click **Run**, type `services.msc`, and click **OK**.
3. In the **Services** window, select the server instance that you want to start, and click **Start**.

What to do next

Because the server service can issue requests that require action, it is important to monitor server activity with the Operations Center or the administrative client.

To view start and stop completion messages that are logged in the Windows application log, use the Event Viewer tool in the **Administrative Tools** folder.

Manually creating and configuring a Windows service

If you configured the server by using the configuration wizard, a Windows service was created automatically. If a service was not created automatically, you must create it.

Before you begin

To complete this procedure, you must log on with a user ID that is in the Administrators group.

Procedure

Open a command window and enter the **sc.exe create** command:

```
sc.exe create server_name binPath= "path_to_server -k instance_name"
start= start_type obj= account_name password= password
```

where:

server_name

Specifies the name of the server service.

path_to_server

Specifies the path to the `dsmsvc.exe` executable file, including the file name. This path is the default path:

`C:\Program Files\Tivoli\TSM\server`

instance_name

Specifies the name of the Db2 instance, which is also the name of the server instance, for example, `Server1`.

start_type

Specifies the method for starting the service. To automatically start the service, enter `auto`. If you specify the `auto` option, the service starts automatically at system startup and restarts automatically whenever the system is restarted. To manually start the service, enter `demand`.

account_name

Specifies the user ID for the account under which the service runs. For example, the account name might be `Administrator`. This parameter is optional. If it is not specified, the Local System account is used.

password

Specifies the password for the *account_name* user account.

Tip: When you enter the command, ensure that you enter a space after each equal sign (=).

Results

The server starts as a Windows service.

Starting the server in the foreground

To directly interact with an IBM Storage Protect server, start the server in the foreground. For example, if you want to enter commands, start the server in the foreground.

Procedure

1. Change to the directory where the server is installed.
For example, change to the `c:\program files\tivoli\tsm\server` directory.
2. Enter the following command:

```
dsmserv -k instance_name
```

where *instance_name* specifies the server instance.

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, `dsmserve.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:

```
dsmserve 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 Storage Protect command line:

```
halt
```

Registering licenses

Immediately register any IBM Storage 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 Storage Protect license.

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

Preparing the server for database backup operations

To prepare the server for automatic and manual database backup operations, ensure that you specify a tape, file, or cloud device class and complete other steps.

Procedure

1. Ensure that the IBM Storage Protect server configuration is complete.

Tip: You can configure the server for database backups by using the configuration wizard (`dsmicfgx`) or you can complete the steps manually. For more information about configuration, see the *Configuring servers* section in IBM Documentation.

2. Select the device class to be used for database backups, protect the master encryption key, and set a password.

Ensure that the following key files are protected:

- Master encryption key files (`dsmkeydb.*`)
- Server certificate and private key files (`cert.*`)

To complete these actions, issue the **SET DBRECOVERY** command from the administrative command line:

```
set dbrecovery device_class_name protectkeys=yes password=password_name
```

where *device_class_name* specifies the device class to be used for database backup operations, and *password_name* specifies the password.

You must specify a device class name or the backup fails. By specifying **PROTECTKEYS=YES**, you ensure that the master encryption key is backed up during database backup operations. Cloud device classes require the **PROTECTKEYS=YES** parameter.

Create a strong password that is at least 8 characters long. If you specify a password for database backup, you must specify the same password on the **RESTORE DB** command to restore the database.



Attention: Ensure that you remember the password and keep a copy stored in a secure location. Without the password, data cannot be recovered.

Example

To specify that database backups include a copy of the master encryption key for the server, run the following command:

```
set dbrecovery dbback protectkeys=yes password=protect8991
```

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 the Creating the server instance section for each new instance, optionally creating 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 V7.1 to V8.1. See the upgrade section 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.

A typical IBM Storage Protect installation involves one server instance on the IBM Storage Protect server computer. You might want to install a second instance if you are configuring in a clustered environment. You might also want to run more than one server on a large computer if you have multiple tape libraries or a disk-only configuration. After you install and configure the first IBM Storage Protect server, use the Server Initialization wizard to create additional IBM Storage Protect server instances on the same computer.

By using the Server Initialization wizard, you can install up to four IBM Storage Protect server instances on a single system or cluster.

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 Storage Protect server fix pack

IBM Storage 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 Storage 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 and then revert the server to an earlier level, 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.

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

Ensure that you retain the installation media from the base release of the installed server. If you installed IBM Storage 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 Storage 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 Storage 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 Storage Protect administrative command:

```
backup db type=dbsnapshot devclass=tapeclass
```

2. Back up the device configuration information. Issue the following IBM Storage 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 Storage 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 Storage Protect installation. The IBM Storage 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. 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](#).
8. Change to the directory where you placed the executable file.

Then, either double-click the following executable file or enter the following command on the command line to extract the installation files.

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.

```
8.x.x.x-IBM-SPSRV-platform.exe
```

where: *platform* denotes the operating system that IBM Storage Protect is to be installed on.

9. Select one of the following ways of installing IBM Storage 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 Storage Protect software by using one of the following methods:

Installation wizard

Follow the instructions for your operating system:

[“Installing IBM Storage Protect by using the installation wizard” on page 64](#)

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 Storage Protect by using console mode” on page 65](#)

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:

C:\ProgramData\IBM\Installation Manager\logs

Applying a fix pack to IBM Storage Protect V8.1 in a clustered environment

IBM Storage 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.

Procedure

1. Complete steps 2 - 6 for each server instance.

2. Back up the database by using the **BACKUP DB** command.

The preferred method is to use a snapshot backup, which provides a full database backup without interrupting scheduled backups.

For example, you can run the following command to create a snapshot backup:

```
backup db type=dbsnapshot devclass=tapeclass
```

3. Back up the device configuration information to another directory by issuing the following command:

```
backup devconfig filenames=file_name
```

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

4. Back up the volume history file to another directory by issuing the following command:

```
backup volhistory filenames=file_name
```

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

5. Save a copy of the server options file, typically named `dsmerv.opt`, which is in the server instance directory.
6. Ensure that the resource group is on the primary node, and that all nodes in the cluster are running. Take the following actions on the primary node:

- a) Take the server resource offline and remove it by completing the following steps:

- i) In the **Failover Cluster Manager** window, select the cluster. Click **Roles** and select the cluster role from the **Name** column. At the end of the **Roles** window, click the **Resources** tab. The server resource is displayed in the **Other Resources** section.
- ii) In the **Other Resources** section, right-click on each entry and view the properties. Review the **Name** and **Type** properties to determine which resource is the IBM Storage Protect server resource and which resource is the Db2 resource.

For example, the IBM Storage Protect server resource might have the following properties:

```
Name: IBM Storage Protect SERVER1 Group Server
Type: TSM Server
```

The Db2 resource might have the following properties:

```
Name: Server1
Type: DB2 Server
```



Warning: Do not remove the Db2 resource. If you unintentionally remove the Db2 resource during this step, you will be unable to recluster after you upgrade the server.

- iii) Select the IBM Storage Protect server resource, and right-click **Take Offline**.
- iv) To remove the IBM Storage Protect server resource, select it, and right-click **Remove**.
- b) In the **Failover Cluster Manager** window, remove the network name and IP address by completing the following steps:
 - i) In the **Server Name** section, expand the network name to view the IP address. Note the network name and IP address. Right-click **Take Offline**.

Note: This action also takes the Db2 server resource offline.

- ii) Select the network name and the IP address, and right-click **Remove**. In the confirmation message, confirm that you want to remove the client access point.
- c) Using the command-line, remove Db2 clustering for each IBM Storage Protect instance in the cluster by issuing the following command for *each* instance:

```
db2mscs -u:instancename
```

For example, to remove Db2 clustering from the SERVER1 instance, issue the following command:

```
db2mscs -u:server1
```

Tip: You might see an error message about a missing cluster resource. Ignore this message.

Note: This action will remove the Db2 resource from the cluster and will delete the DB2 -DB2TSM1 – SERVERX service from the secondary node.

- d) On the primary node, in the **Failover Cluster Manager** window, review the resource group **Summary** section. Verify that only the shared disks and any tape resources remain in the resource group.
7. Stop the cluster service on all nodes in the cluster. One way to stop it is by using the Services Application. Right-click **Cluster Service** and select **Stop**.
 8. Beginning with the primary node, install the IBM Storage Protect V8.1 server on all nodes in the cluster. For instructions, see [Chapter 2, “Installing the server components,” on page 63](#). To install the server by using the installation wizard, click the **Update** icon on the **IBM Installation Manager** window. Do not click the **Install** or **Modify** icons.
 9. Restart the cluster service on all nodes in the cluster. One way to start it is by using the Services Application. Right-click **Cluster Service** and select **Start**. Verify that all cluster resources are on the primary node.
 10. With cluster resources on the primary node, start the configuration wizard by clicking **Start > Apps > IBM Storage Protect > Configuration Wizard**. Complete the following steps in the configuration wizard:
 - a) In the **Instance User ID** panel, enter the following:
 - The Instance name that you want to cluster.
 - The User ID that is associated with the cluster. Domain accounts use the following format: domain\account_name
 - The account password.
 - b) When prompted, click **Yes** to confirm that you want to recluster.
 - c) Continue all steps in the wizard until you see a message confirming that the configuration was successful.

Note: In the Cluster Manager, the status for the Db2 and IBM Storage Protect Group Server resources might be in a failed state until the server finishes retrying startup.

- d) If your environment has multiple server instances, complete step 9 for each server instance.
- 11. Optional: If necessary, register licenses for the IBM Storage Protect server components that are installed on your system by issuing the **REGISTER LICENSE** command:

```
register license file=installation_directory\server\component_name.lic
```


Chapter 5. Upgrading to V8.1

To take advantage of new product features and updates, upgrade the IBM Storage Protect server.

Before you begin

Review the security updates planning information in [“What you should know about security before you install or upgrade the server”](#) on page 3.

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 Storage Protect Upgrade and Migration Process - Frequently Asked Questions](#).

Table 13. 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 Storage Protect server fix pack,” on page 89
V7.1	V8.1	“Installing the server and verifying the upgrade” on page 99
V5.5, V6.2, or V6.3	V8.1	IBM Storage Protect Upgrade and Migration Process - Frequently Asked Questions

An upgrade 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 102.

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 Storage 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” on page 96](#)
2. [“Preparing the system” on page 96](#)
3. [“Installing the server and verifying the upgrade” on page 99](#)

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:

[System requirements for Windows](#)

For the latest updates related to system requirements, see the IBM Storage 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_8.1.11/srv.common/r_relnotes_srv.html) and readme files for server components.
3. Review the security updates planning information in [“What you should know about security before you install or upgrade the server” on page 3](#).
4. 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.
5. If you are upgrading the server from V7 to V8.1, verify that you have the system ID and password for the IBM Db2 instance of the IBM Storage 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 IBM 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 administrative user ID that was used to install the V7.1 server.

2. Obtain a list of Db2 instances. Issue the following system command:

```
db2ilist
```

The output might be similar to the following example:

```
SERVER1
```

Ensure that each instance corresponds to a server that is running on the system.

3. Gather information about 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.

- a) Open the Db2 command window by issuing the following system command:

```
db2cmd
```

- b) To change the instance, issue the following system command:

```
set DB2INSTANCE=instance
```

where *instance* specifies the Db2 instance.

- c) Obtain the default database path for the Db2 instance by issuing the following system command:

```
db2 get dbm cfg | findstr DFTDBPATH
```

The output might be similar to the following example:

```
Default database path                (DFTDBPATH) = D:
```

- d) Obtain information about the Db2 instance databases by issuing the following system command:

```
db2 list database directory
```

The output might be similar to the following example:

```
System Database Directory
    Number of entries in the directory = 2
    Database 1 entry:
        Database alias                = TSMAL001
        Database name                 = TSMDB1
        Node name                     = TSMNODE1
        Database release level        = d.00
        Comment                       = TSM SERVER DATABASE VIA TCP/IP
        Directory entry type          = Remote
        Catalog database partition number = -1
        Alternate server hostname      =
        Alternate server port number   =
    Database 2 entry:
        Database alias                = TSMDB1
        Database name                 = TSMDB1
        Local database directory       = D:
        Database release level        = d.00
        Comment                       =
        Directory entry type          = Indirect
        Catalog database partition number = 0
        Alternate server hostname      =
        Alternate server port number   =
```

- e) Obtain the Db2 instance variables by issuing the following system command:

```
db2set -all
```

The output might be similar to the following example:

```
[e] DB2CODEPAGE=1208
[e] DB2PATH=D:\TSM\db2
[i] DB2_PMODEL_SETTINGS=MAX_BACKGROUND_SYSAPPS:500
[i] DB2_SKIPINSERTED=ON
```

```
[i] DB2_KEEPTABLELOCK=OFF
[i] DB2_EVALUNCOMMITTED=ON
[i] DB2_VENDOR_INI=D:\Server1\tsmdbmgr.env
[i] DB2_SKIPDELETED=ON
[i] DB2INSTPROF=C:\ProgramData\IBM\DB2\DB2TSM1
[i] DB2COMM=TCPIP
[i] DB2CODEPAGE=819
[i] DB2_PARALLEL_IO=*
[g] DB2_EXTSECURITY=YES
[g] DB2_COMMON_APP_DATA_PATH=C:\ProgramData

[g] DB2PATH=D:\TSM\db2
[g] DB2INSTDEF=SERVER1
```

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 `dsmsevr.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.

Open the Windows Task Manager application and review the list of active processes.

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 with the administrative user ID that was used to install the previous server.

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

About this task

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

- Server

Tip: The database (IBM 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 Storage 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:

- a. 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 Storage Protect [technote 588021](#)
- IBM Storage Protect Extended Edition [technote 588023](#)
- IBM Storage Protect for Data Retention [technote 588025](#)

- b. Change to the directory where you placed the executable file.

Tip: In the next step, the files are extracted to the current directory. The path must contain no more than 128 characters. 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. To extract the installation files, double-click the executable file:

```
package_name.exe
```

Where *package_name* is similar to the following example:

```
8.1.x.000-SPSRV-WindowsX64.exe
```

The package is large. Therefore, the extraction takes some time.

3. Install the IBM Storage Protect software by using one of the following methods. Install the IBM Storage Protect license during the installation process.

Tip: If you have multiple server instances on your system, install the IBM Storage 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 Storage Protect by using the installation wizard”](#) on page 64.

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 Storage Protect by using console mode”](#) on page 65.

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 Storage Protect in silent mode”](#) on page 67.

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:

```
C:\ProgramData\IBM\Installation Manager\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.
- 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 Storage Protect administrative command:

```
dsmdmc
```

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

```
query system
```

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

```
query db format=detailed
```

7. Verify that the upgrade was successful:

- a) Start the server instance. To start the server from the default directory, C:\Program Files\Tivoli\TSM, run the following IBM Storage Protect administrative command:

```
dsmserve -k server_instance
```

server_instance is the name of your server instance. Server1 is the default name for the first instance of the IBM Storage Protect server.

If you plan to run the server as a service under the Local System account, the Local System account must be explicitly granted access to the server database. For instructions, see [“Starting the server as a Windows service”](#) on page 82.

- 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 Storage Protect administrative command:

```
dsmadm
```

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

```
query system
```

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

```
query db format=detailed
```

8. Register the licenses for the server components that are installed on your system by running the **REGISTER LICENSE** command:

```
register license file=installation_directory\server\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, c:\Program Files\Tivoli\TSM, run the following command to register the license:

```
register license file=c:\Program Files\Tivoli\TSM\server\tsmbasic.lic
```

For example, if you installed IBM Storage Protect Extended Edition in the c:\Program Files\Tivoli\TSM directory, run the following command:

```
register license file=c:\Program Files\Tivoli\TSM\server\tsmee.lic
```

For example, if you installed IBM Storage Protect for Data Retention in the c:\Program Files\Tivoli\TSM directory, run the following command:

```
register license file=c:\Program Files\Tivoli\TSM\server\dataret.lic
```

Restriction:

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

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

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

9. Prepare the server for automatic and manual database backup operations.

For instructions, see [“Preparing the server for database backup operations”](#) on page 85.

10. Verify that the PASSWORDACCESS GENERATE application programming interface (API) client option is not in the client options file. The default name for the client options file is `dsm.opt` and it is typically located in the `c:\programdata\ibm\db2\db2tsm1\tsminst1` directory. Remove the PASSWORDACCESS GENERATE option if it is present.
11. Optional: To install an extra language package, use the modify function of the IBM Installation Manager.
12. Optional: To upgrade to a newer version of a language package, use the update function of the IBM Installation Manager.
13. To facilitate troubleshooting in case of any future issues, ensure that sufficient space is allocated for a core dump. For more information, see [technote 6357399](#).

What to do next

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

If a device driver is available on Windows for the tape drives or medium changers that you plan to use, use the device driver. If a device driver is not available on Windows for the tape drives or medium changers that you plan to use, install the IBM Storage Protect device driver by running the `dpinst.exe /a` command. The `dpinst.exe` file is in the device driver directory. The default directory is `C:\Program Files\Tivoli\TSM\device\drivers`.

Upgrading the server in a clustered environment

To upgrade a server 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 14. Procedures for upgrading the server in a clustered environment on a Windows operating system		
Source release	Target release	Procedure
V8.1	V8.1 fix pack	“Applying a fix pack to IBM Storage Protect V8.1 in a clustered environment” on page 91
V6.3 or V7.1	V8.1	Upgrading V6.3 or V7.1 to V8.1 in a clustered environment on Windows
V5.5, V6.1, V6.2	V7.1 or later	IBM Storage Protect Upgrade and Migration Process - Frequently Asked Questions

Upgrading a V7.1 server to V8.1 in a clustered environment

To take advantage of new product features, you can upgrade a server that is installed on a Windows operating system in a clustered environment.

Before you begin

Ensure that you retain the installation media from the V7.1 server base release that you are upgrading. If you installed the server from a downloaded package, ensure that the downloaded files are available.

If the upgrade fails, and the server license module is uninstalled, you must reinstall the license from the installation media of the server base release.

Procedure

1. Complete steps 2 - 6 for each server instance.

2. Back up the database by using the **BACKUP DB** command.

The preferred method is to use a snapshot backup, which provides a full database backup without interrupting scheduled backups.

For example, you can run the following command to create a snapshot backup:

```
backup db type=dbsnapshot devclass=tapeclass
```

3. Back up the device configuration information to another directory by issuing the following command:

```
backup devconfig filenames=file_name
```

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

4. Back up the volume history file to another directory by issuing the following command:

```
backup volhistory filenames=file_name
```

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

5. Save a copy of the server options file, typically named `dsmserv.opt`, which is in the server instance directory.
6. Ensure that the resource group is on the primary node, and that all nodes in the cluster are running. Take the following actions on the primary node:

- a) Take the server resource offline and remove it by completing the following steps:

- i) In the **Failover Cluster Manager** window, select the cluster. Click **Roles** and select the cluster role from the **Name** column. At the end of the **Roles** window, click the **Resources** tab. The server resource is displayed in the **Other Resources** section.
- ii) In the **Other Resources** section, right-click on each entry and view the properties. Review the **Name** and **Type** properties to determine which resource is the IBM Storage Protect server resource and which resource is the Db2 resource.

For example, the IBM Storage Protect server resource might have the following properties:

```
Name: IBM Storage Protect SERVER1 Group Server
Type: TSM Server
```

The Db2 resource might have the following properties:

```
Name: Server1
Type: DB2 Server
```



Warning: Do not remove the Db2 resource. If you unintentionally remove the Db2 resource during this step, you will be unable to recluster after you upgrade the server.

- iii) Select the IBM Storage Protect server resource, and right-click **Take Offline**.
 - iv) To remove the IBM Storage Protect server resource, select it, and right-click **Remove**.
- b) In the **Failover Cluster Manager** window, remove the network name and IP address by completing the following steps:
 - i) In the **Server Name** section, expand the network name to view the IP address. Note the network name and IP address. Right-click **Take Offline**.

Note: This action also takes the Db2 server resource offline.

 - ii) Select the network name and the IP address, and right-click **Remove**. In the confirmation message, confirm that you want to remove the client access point.

- c) Using the command line, remove Db2 clustering for each IBM Storage Protect instance in the cluster by issuing the following command for *each* instance:

```
db2mscs -u:instancename
```

For example, to remove Db2 clustering from the SERVER1 instance, issue the following command:

```
db2mscs -u:server1
```

Tip: You might see an error message about a missing cluster resource. Ignore this message.

Note: This action removes the Db2 resource from the cluster and deletes the **DB2 -DB2TSM1 – SERVERX** service from the secondary node.

- d) On the primary node, in the **Failover Cluster Manager** window, review the resource group **Summary** section. Verify that only the shared disks and any tape resources remain in the resource group.
7. Stop the cluster service on all nodes in the cluster. One way to stop it is by using the Services Application. Right-click **Cluster Service** and select **Stop**.
8. Beginning with the primary node, install the IBM Storage Protect V8.1 server on all nodes in the cluster. For instructions, see Chapter 2, “Installing the server components,” on page 63. To install the server by using the installation wizard, click the **Update** icon on the **IBM Installation Manager** window. Do not click the **Install** or **Modify** icons.
9. With cluster resources on the primary node, start the configuration wizard by clicking **Start > Apps > IBM Storage Protect > Configuration Wizard**. Complete the following steps in the configuration wizard:
- a) In the **Instance User ID** window, enter the following information:
- The Instance name that you want to cluster.
 - The User ID that is associated with the cluster. Domain accounts use the following format: domain\account_name
 - The account password.
- b) When prompted, click **Yes** to confirm that you want to recluster.
- c) Continue all steps in the wizard until you see a message that confirms that the configuration was successful.
- Note:** In the Cluster Manager, the status for the Db2 and IBM Storage Protect Group Server resources might be in a failed state until the server finishes retrying startup.
- d) Set up services on all secondary nodes by issuing the following command for each secondary node:

```
C:\windows\system\sc.exe create "IBM Storage Protect instance_name"  
binPath= "server_installation_directory\dsmsvc.exe -k instance_name"  
start = demand obj=instance_username password = instance_user's password
```

For example, issue the following command:

```
C:\Windows\system32\sc.exe create "IBM Storage Protect SERVER1"  
binPath= "C:\TSM\Server\dsmsvc.exe -k SERVER1" start= demand  
obj= tsm13\Administrator password= "PassW0rd"
```

- e) If your environment has multiple server instances, complete step 9 for each server instance.
10. Optional: If necessary, register licenses for the IBM Storage Protect server components that are installed on your system by issuing the **REGISTER LICENSE** command:

```
register license file=installation_directory\server\component_name.lic
```

Upgrading IBM Storage Protect servers in a clustered HADR environment

You can upgrade both the primary and standby servers in a *high availability disaster recovery* (HADR) environment.

Procedure

1. On the primary server, stop the IBM Storage Protect server instance.
2. On the standby server, stop the IBM Storage Protect server instance and the HADR IBM Db2 instance.
3. On the primary server, issue the following commands:

```
db2start  
db2ckupgrade tsmbd1 -l /tmp/upgrade.out  
db2stop
```

4. Upgrade the primary IBM Storage Protect server. For instructions, see [“Upgrading the server in a clustered environment” on page 102](#).
5. On the standby server, issue the following commands:

```
db2start  
db2ckupgrade tsmbd1 -l /tmp/upgrade.out  
db2stop
```

6. Upgrade the standby IBM Storage Protect server. For instructions, see [“Upgrading the server in a clustered environment” on page 102](#).
7. On the standby server, start HADR Db2.
8. On the primary server, start Db2 and the IBM Storage Protect server instance.

Chapter 6. Reverting to the previous server version in a cluster configuration

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 Storage Protect server with minimal loss of data.

Before you begin

If you must revert to the previous version of the server after an upgrade, you must have the server installation media for your original version. You must also have the following items from the earlier version of the server:

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

Steps for reverting to the previous server version

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 server with minimal loss of data.

About this task

Complete the following steps on the system that you are preparing, for example, the Version 8.1.7 server.

Procedure

1. Delete the IBM Storage Protect server resource and the network resource in Microsoft Failover Cluster Manager.
 - a. Open Failover Cluster Manager. Under **Other Resources**, right-click the IBM Storage Protect instance resource. Select **Take this resource offline**.
 - b. Delete the IBM Storage Protect instance.
 - c. Under **Server Name**, expand the network name and record the cluster TCP/IP address.
 - d. Delete the network name.
2. Remove the IBM Db2 cluster from the instance by issuing the following command:

```
DB2mcs -u:instance_name
```

For example, if the server instance is Server1, enter the command:

```
db2mcs -u:Server1
```

Tip: You might see an error message about a missing cluster resource. Disregard this message.

3. Remove the database.

One way to remove it is by issuing this command:

```
dsmserve -k instance_name removedb tsmdb1
```

4. On each system in the cluster, delete the V 8.1.7 tsmsvrsc DLL files by completing the following steps:
 - a. Stop the cluster service. One way to stop it is by using the Services Application. Right-click **Cluster Service** and select **Stop**.

- b. Delete the `tsmsvrrscexX64.dll` and `tsmsvrrscx64.dll` files from the `C:\Windows\Cluster` directory.
 - c. Start the cluster service. One way to start it is by using the Services Application. Right-click **Cluster Service** and select **Start**.
5. Use the uninstallation program to uninstall the V8.1 server. Uninstallation removes the server and the database manager, with their directories.
6. Clean up the database and recovery log directories if you are reusing them.
7. Stop the cluster service. Reinstall the version of the server program that you were using before the upgrade to V8.1.7. 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 V7.1.8 before the upgrade, and you intend to use the database backup that was in use on this server. You must install the V7.1.8 fix pack to be able to restore the database backup.
8. Copy the following files to the instance directory.
 - Device configuration file
 - Volume history file
 - The server options file (typically `dsmserv.opt`)
9. Use the configuration wizard (`dsmicfgx`) to recreate the server instance.
10. Restore the database to a point in time before the upgrade.

Chapter 7. Reference: IBM Db2 commands for IBM Storage 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 Storage 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 the table.

This list is supplemental material only and is not a comprehensive list. There is no implication that an IBM Storage 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 documentation.

Table 15. Db2 commands		
Command	Description	Example
db2cmd	Opens the Db2 command line processor window and initializes the command-line environment.	Open the Db2 command window: db2cmd
db2icrt	Creates Db2 instances in the home directory of the instance owner. Tip: The IBM Storage 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 located in the DB2PATH\bin directory where DB2PATH is the location where the Db2 copy is installed.	Manually create an IBM Storage 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 Storage Protect database from a Db2 CLI: db2 connect to tsmdb1

Table 15. Db2 commands (continued)

Command	Description	Example
GET DATABASE CONFIGURATION	<p>Returns the values of individual entries in a specific database configuration file.</p> <p>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 Storage Protect server commands or procedures.</p>	<p>Show the configuration information for a database alias:</p> <pre>db2 get db cfg for tsmdb1</pre> <p>Retrieve information in order to verify settings such as database configuration, log mode, and maintenance.</p> <pre>db2 get db config for tsmdb1 show detail</pre>
GET DATABASE MANAGER CONFIGURATION	<p>Returns the values of individual entries in a specific database configuration file.</p> <p>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 Storage Protect server commands or procedures.</p>	<p>Retrieve configuration information for the database manager:</p> <pre>db2 get dbm cfg</pre>
GET HEALTH SNAPSHOT	<p>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.</p> <p>IBM Storage 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 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.</p> <p>IBM Storage Protect monitors the condition and responds appropriately. Not all declared alerts by the Db2 database are acted on.</p>	<p>Receive a report on Db2 health monitor indicators:</p> <pre>db2 get health snapshot for database on tsmdb1</pre>
GRANT (Database Authorities)	<p>Grants authorities that apply to the entire database rather than privileges that apply to specific objects within the database.</p>	<p>Grant access to the user ID itmuser:</p> <pre>db2 GRANT CONNECT ON DATABASE TO USER itmuser db2 GRANT CREATETAB ON DATABASE TO USER itmuser</pre>

Table 15. Db2 commands (continued)		
Command	Description	Example
RUNSTATS	<p>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.</p> <p>To see a table, issue this utility after updating or reorganizing the table.</p> <p>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.</p> <p>Tip: The server configures Db2 to run the RUNSTATS command as needed.</p>	<p>Update statistics on a single table.</p> <pre>db2 runstats on table SCHEMA_NAME.TABLE_NAME with distribution and sampled detailed indexes all</pre>
set db2instance	Determines which instance applies to the current session.	<p>Determine which instance is applicable:</p> <pre>set db2instance=tsminst1</pre>
SET SCHEMA	<p>Changes the value of the CURRENT SCHEMA special register, in preparation for issuing SQL commands directly through the Db2 CLI.</p> <p>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.</p>	<p>Set the schema for IBM Storage Protect:</p> <pre>db2 set schema tsmdb1</pre>
START DATABASE MANAGER	<p>Starts the current database manager instance background processes. The server starts and stops the instance and database whenever the server starts and halts.</p> <p>Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.</p>	<p>Start the database manager:</p> <pre>db2start</pre>
STOP DATABASE MANAGER	<p>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.</p> <p>This command is not valid on a client.</p> <p>The server starts and stops the instance and database whenever the server starts and halts.</p> <p>Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.</p>	<p>Stop the database manager:</p> <pre>db2 stop dbm</pre>

Chapter 8. Uninstalling IBM Storage Protect

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

Before you begin

Complete the following steps before you uninstall IBM Storage 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.



Attention: Do not use the Add/Remove Programs tool in the Windows Control Panel to uninstall IBM Storage Protect. Use only the uninstallation procedure that is described in this section.

About this task

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

What to do next

Reinstall the IBM Storage Protect components.

Uninstalling IBM Storage Protect by using a graphical wizard

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

Procedure

1. Start the Installation Manager.
Open the Installation Manager from the **Start** menu.
2. Click **Uninstall**.
3. Select **IBM Storage Protect server**, and click **Next**.
4. Click **Uninstall**.
5. Click **Finish**.

Uninstalling IBM Storage Protect in console mode

To uninstall IBM Storage 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:
`C:\Program Files\IBM\Installation Manager\eclipse\tools`
2. From the `tools` directory, issue the following command:
`imcl.exe -c`
3. To uninstall, enter 5.

4. Choose to uninstall from the IBM Storage Protect package group.
5. Enter N for Next.
6. Choose to uninstall the IBM Storage Protect server package.
7. Enter N for Next.
8. Enter U for Uninstall.
9. Enter F for Finish.

Uninstalling IBM Storage Protect in silent mode

To uninstall IBM Storage 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 Storage Protect server components. IBM Storage 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 Storage 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:

```
C:\Program Files\IBM\Installation Manager\eclipse\tools
```

2. From the tools directory, issue the following command, where *response_file* represents the response file path, including the file name:

```
imcl.exe -input response_file -silent
```

The following command is an example:

```
imcl.exe -input C:\tmp\input\uninstall_response.xml -silent
```

Uninstalling and reinstalling IBM Storage Protect

If you plan to manually reinstall IBM Storage 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 Storage Protect, complete the following steps:

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

```
db2ilist
```

2. Run the following commands for every server instance:

```
db2 attach to server1
db2 get dbm cfg show detail
db2 detach
```

Keep a record of the database path for each instance.

3. Uninstall IBM Storage Protect.

After uninstalling IBM Storage Protect, check the **Control Panel > Add or Remove Programs** to verify that IBM Storage Protect IBM Db2 is uninstalled.

4. When you uninstall any supported version of IBM Storage Protect, including a fix pack, an instance file is created. The instance file is created to help reinstall IBM Storage 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:

`C:\ProgramData\IBM\Tivoli\TSM\instanceList.obj` in the IBM Storage Protect server installation directory

5. Reinstall IBM Storage Protect.

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

- a. Recreate your server instances.

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:

```
set db2instance=server1
db2 catalog database tsmdb1
db2 attach to server1
db2 update dbm cfg using dftdbpath instance_drive
db2 detach
```

- c. Verify that IBM Storage 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, click **Start > All Programs > IBM Installation Manager > View Installed Packages**.

Procedure

To uninstall IBM Installation Manager, complete the following steps:

Uninstalling IBM Storage Protect

- 1. From the **Start** menu, click **Control Panel > Programs and Features**.
 2. Select **IBM Installation Manager** and click **Uninstall**.

Part 2. Installing and upgrading the Operations Center

The IBM Storage 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](#)
 - [Operations Center computer requirements](#)
 - [Hub and spoke server requirements](#)
 - [Operating system requirements](#)
 - [Web browser requirements](#)
 - [Language requirements](#)
 - [Requirements and limitations for IBM Storage Protect client management services](#)
- [Administrator IDs that the Operations Center requires](#)
- [IBM Installation Manager](#)
- [Installation checklist](#)
- [Obtaining the Operations Center installation package](#)

About this task

Table 16 on page 117 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 [Upgrading the Operations Center](#).

Table 16. Methods for installing or uninstalling the Operations Center	
Method	Instructions
Graphical wizard	<ul style="list-style-type: none">• Installing the Operations Center by using a graphical wizard• Uninstalling the Operations Center by using a graphical wizard
Console mode	<ul style="list-style-type: none">• Installing the Operations Center in console mode• Uninstalling the Operations Center in console mode
Silent mode	<ul style="list-style-type: none">• Installing the Operations Center in silent mode• “Uninstalling the Operations Center in silent mode” on page 183

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 Storage 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 Storage 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 17 on page 119 lists the prerequisite requirements that are verified during the installation and indicates where to find more information about these requirements.

Table 17. Requirements that are verified during the installation	
Requirement	Details
Minimum memory requirement	“Operations Center computer requirements” on page 120
Operating system requirement	“Operating system requirements” on page 122
Host name for the computer where the Operations Center will be installed	“Installation checklist” on page 126
Requirements for the Operations Center installation directory	“Installation checklist” on page 126

Operations Center computer requirements

You can install the Operations Center on a computer that is also running IBM Storage 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

For the most up-to-date requirements information, see [Software and Hardware Requirements](#).

The hub and spoke servers that are monitored by the Operations Center require additional resources, as described in [“Hub and spoke server requirements”](#) on page 120.

Hub and spoke server requirements

When you open the Operations Center for the first time, you must associate the Operations Center with one IBM Storage 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 18 on page 120 indicates the version of IBM Storage Protect server that must be installed on the hub server and on each spoke server that is managed by the Operations Center.

Table 18. IBM Storage Protect server version requirements for hub and spoke servers		
Operations Center	Version on the hub server	Version on each spoke server
8.1.14	8.1.14	8.1.12 or later Restrictions: <ul style="list-style-type: none"> Some Operations Center functions are not available for servers that use a version earlier than 8.1.14. A spoke server cannot use a version that is later than the version on the hub server.

For information about hub and spoke server compatibility requirements for other versions of the Operations Center, see [technote 496593](#).

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 Storage Protect on each spoke server. However, a general guideline is that a hub server on a separate system, such as a VM, can support dozens of 8.1 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 Storage 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 spoke servers than can be managed on one hub, 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.

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. 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.

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 the specific browser information and supported versions, see [Software and Hardware Requirements](#).

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.

Communication between the Operations Center and the web browser must be secured by using the Transport Layer Security (TLS) 1.2 protocol. The web browser must support TLS 1.2, and TLS 1.2 must be enabled. The web browser displays an SSL error if it does not meet these 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.

<i>Table 19. Operations Center language values that you can use on Windows systems</i>	
Language	Language option value
Chinese, Simplified	chs
Chinese, Traditional	cht
English	ameng
French	fra
German	deu
Italian	ita
Japanese (Shift-JIS)	jpn
Korean	kor
Portuguese, Brazilian	ptb
Russian	rus
Spanish	esp

Requirements and limitations for IBM Storage Protect client management services

For versions of IBM Storage Protect that are earlier than Version 8.1.13, the client management service is a separate component that you install on backup-archive clients. This component collects diagnostic information such as client log files. Before you install the client management service on your system, you must understand the requirements and limitations.

Restriction: The following requirements apply only to client versions that are earlier than 8.1.13. In IBM Storage Protect 8.1.13, the installation of a separate client management service package was deprecated and the feature that the client management service provided was integrated into the backup-archive client package.

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.

Tip: Before you install the client management service, ensure that a successful connection was established between the backup-archive client and the server. The server truststore file that the client uses does not have the server Secure Sockets Layer (SSL) certificate until the client system has connected to the server.

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) version 1.2 or later must be installed for transmission of data between the client management service and the Operations Center. Basic authentication is provided and data and authentication information are encrypted through the Secure Sockets Layer (SSL) channel. TLS is automatically installed along with the necessary SSL certificates when you install the client management service.

Beginning with IBM Storage Protect Version 8.1.11, the TLS 1.3 protocol is enabled by default to secure communications between servers, clients, and storage agents. To use TLS 1.3, both parties in the communication session must use TLS 1.3. If either party uses TLS 1.2, then both parties use TLS 1.2 by default.
- 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 Storage Protect for Virtual Environments: Data Protection for VMware.
- You cannot install the client management service on other IBM Storage 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 Storage Protect servers that are installed on the same server, you can view log files for only one of the client nodes.

To learn about possible updates related to the client management service, see [technote 534165](#).

Related tasks

“Collecting diagnostic information with IBM Storage Protect client management services” on page 161
The client management service collects diagnostic information about backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

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 Storage 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 120.

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 20 on page 126](#) 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.

Table 20. 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: <ul style="list-style-type: none">• 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: <ul style="list-style-type: none">• Administrator

Table 20. Information to verify or determine before you install the Operations Center (continued)	
Information	Details
Operations Center installation directory	<p>The Operations Center is installed in the <code>ui</code> subdirectory of the installation directory.</p> <p>The following path is the default path for the Operations Center installation directory:</p> <ul style="list-style-type: none"> <code>c:\Program Files\Tivoli\TSM</code> <p>For example, if you use this default path, the Operations Center is installed in the following directory:</p> <pre>c:\Program Files\Tivoli\TSM\ui</pre> <p>The installation directory path must meet the following criteria:</p> <ul style="list-style-type: none"> 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: <pre>% < > ' " \$ & ; *</pre>
IBM Installation Manager installation directory	<p>The following path is the default path for the IBM Installation Manager installation directory:</p> <ul style="list-style-type: none"> <code>C:\Program Files\IBM\Installation Manager</code>
Port number that is used by the Operations Center web server.	<p>The value for the secure (https) port number must meet the following criteria:</p> <ul style="list-style-type: none"> The number must be an integer in the range 1024 - 65535. The number cannot be in use or allocated to other programs. <p>If you do not specify a port number, the default value is 11090.</p> <p>Tips:</p> <ul style="list-style-type: none"> Although you must specify an integer in the range 1024 - 65535, you can later configure the Operations Center to use the standard TCP/IP secure port (port 443). For more information, see “Configuring the Operations Center web server to use the standard TCP/IP secure port” on page 140. If you later do not remember the port number that you specified, refer to the following file, where <code>installation_dir</code> represents the directory where the Operations Center is installed: <ul style="list-style-type: none"> <code>installation_dir\ui\Liberty\usr\servers\guiServer\bootstrap.properties</code> <p>The <code>bootstrap.properties</code> file contains the IBM Storage Protect server connection information.</p>

Table 20. Information to verify or determine before you install the Operations Center (continued)

Information	Details
Password for secure communications	<p>The Operations Center uses Hypertext Transfer Protocol Secure (HTTPS) to communicate with web browsers.</p> <p>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.</p> <p>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.</p> <p>The password for the truststore file must meet the following criteria:</p> <ul style="list-style-type: none"> • The password must contain a minimum of 6 characters and a maximum of 64 characters. • The password must contain at least the following characters: <ul style="list-style-type: none"> – 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: <div data-bbox="553 1045 1468 1100" style="background-color: #f0f0f0; padding: 5px;"> ~ @ # \$ % ^ & * _ - + = ` </div> <div data-bbox="553 1113 1468 1167" style="background-color: #f0f0f0; padding: 5px;"> () { } [] : ; < > , . ? / </div>

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 Storage 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 120.

You can install the Operations Center on a computer with the IBM Storage 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.

- a. Download the following package file to the directory of your choice:

```
version_number.000-IBM-SPOC-WindowsX64.exe
```

- b. In Windows Explorer, double-click the file name to extract the installation files.

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.bat
```
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 135.

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.bat -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 135](#).

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 126](#).

To encrypt the password, follow the instructions in [“Encrypting passwords in silent installation response files” on page 131](#).

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.bat -s -input response_file -acceptLicense`

What to do next

See [“Configuring the Operations Center”](#) on page 135.

Encrypting passwords in silent installation response files

For added security during a silent installation of the Operations Center, you can encrypt the password in the response file. Only one password (encrypted or unencrypted) can be listed in the data key field in the response file.

Before you begin

Open IBM Installation Manager. In the directory where IBM Installation Manager is installed, go to the `eclipse` subdirectory. By default, the subdirectory is in the following location:

`C:\Program Files\IBM\Installation Manager\eclipse`

Procedure

To encrypt the password in the response file that is used to silently install the Operations Center and ensure that only one password is used in the data key field, complete the following steps:

1. Issue the following command on one line:

```
IBMIM.exe -silent -noSplash encryptString string_to_encrypt
>encrypted_pwd.txt
```

where *string_to_encrypt* is the value that is encrypted and *encrypted_pwd* is the file that contains the encrypted value.

2. Open the encrypted password file and copy the value into the data key field of the response file. Then, remove the encrypted password file by commenting it out.
3. To remove the non-encrypted password from the data key field, complete the following steps:
 - a. Comment out the non-encrypted password (`user.SSL_PASSWORD`) so that the password row is similar to the following example:

```
<!-- <data key='user.SSL_PASSWORD' value='${ssl.password}' /> -->
```

- b. Remove the comment tags from the encrypted password (`user.SSL_PASSWORD_ENCRYPTED`) so that the password rows are similar to the following example:

```
<data key='user.enableSP800_131' value='${enable.SP800131a}' />
<data key='user.SSL_PASSWORD_ENCRYPTED' value='${ssl.password.encrypted}' />
```

Restriction: Use only one value in the data key field in the response file, either the `user.SSL_PASSWORD` or the `user.SSL_PASSWORD_ENCRYPTED` password. You must comment out the one that you are not using or you will receive an error message and the installation will fail.

Example

Using the Installation Manager command line tool, encrypt the password `passw0rd`. Save the encrypted value to the `my_pwd.txt` file. Issue the following command:

```
IBMIM.exe -silent -noSplash encryptString passw0rd > my_pwd.txt
```

where the `my_pwd.txt` file contains the encrypted value, `rbN1IaMAWYYtQxLf6KdNyA==`:

```
<variable name='ssl.password.encrypted' value=' rbN1IaMAWYYtQxLf6KdNyA==' />
```

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 2 on page 135 illustrates an Operations Center configuration.




Figure 2. 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 Storage Protect server that is

designated as the hub server. You can then connect additional IBM Storage 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 Storage 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 communications between the Operations Center and the hub server by using self-signed certificates” on page 142.](#)

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 Storage 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.

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 Storage 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:

- Suspend notification about individual alerts, as described in [“Suspending email alerts temporarily” on page 139](#).
- 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 137](#).

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.
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
```

Configuring the Operations Center web server to use the standard TCP/IP secure port

Port 443 is the standard port for secure web browser communication. If users must access the Operations Center through a firewall, you can configure the Operations Center to communicate through this standard port. In this way, you can avoid opening another port in the firewall.

About this task

When you install the Operations Center, the default port number for secure communication between the Operations Center web server and web browsers is 11090. You can accept this default port at installation time, or you can specify a different port number in the range 1024 - 65535. You cannot specify a port number that is less than 1024 at installation time because those ports are reserved for specific network services.

After the Operations Center is installed, the web server listens on the specified port for requests from web browsers. If users are unable to open the Operations Center because the port is blocked by a firewall, an administrator must open the port to allow browsers to connect. In some production environments, it might be more efficient to use system port 443. Because this system port is reserved for secure web browsing, it is likely already an open port in the firewall. Although you cannot specify port 443 at installation time, you can specify this port after installation.

Procedure

To configure the Operations Center web server to use port 443, complete the following steps after you install the Operations Center:

1. Stop the Operations Center web server.
For instructions about stopping the web server, see [“Starting and stopping the web server” on page 160](#).
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 a property that specifies the port that the Operations Center web server uses for secure communication.
4. Update the `tsm.https.port` property to specify port 443:

```
tsm.https.port=443
```

5. Save and close the `bootstrap.properties` file.
6. Start the Operations Center web server.

For instructions about starting the Operations Center web server, see [“Starting and stopping the web server” on page 160](#).

What to do next

Notify users that the Operations Center is using the standard TCP/IP secure port. Typically, a user opens the Operations Center in their browser by including the port number in the URL. Because port 443 is the default for secure web browser communication, users do not have to specify the port number in the URL. Instead, the following URL can be used, where *hostname* specifies the name of the computer where the Operations Center is installed:

```
https://hostname/oc/
```

For instructions about opening the Operations Center, see [“Opening the Operations Center” on page 160](#).

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-01.ibm.com/support/docview.wss?uid=swg21997347>, 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 Version 1.2 or later is required for secure communication between the IBM Storage Protect server and the Operations Center, and between the hub server and spoke servers.

Securing communications between the Operations Center and the hub server by using self-signed certificates

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. During the installation of the Operations Center, you must create a password for the truststore file. To secure communications 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 must now re-create and configure the truststore file. For instructions, see "Deleting and reassigning the password for the Operations Center truststore file" in IBM Documentation.

The following figure illustrates the components for setting up a Secure Sockets Layer (SSL) connection between the hub server and the Operations Center.



About this task

This procedure provides steps to implement secure communications by using self-signed certificates. If you use certificates that are signed by a certificate authority (CA), see [Securing communications between the Operations Center and the hub server by using CA-signed certificates](#).

Procedure

1. Stop the Operations Center web server.
2. Go to the command line of the operating system on which the Operations Center is installed.
3. Add the certificate to the truststore file of the Operations Center by using the **iKeycmd** utility or the **iKeyman** utility.

The **iKeycmd** utility is a command-line interface, and the **iKeyman** utility is the IBM Key Management graphical user interface.

The **iKeycmd** and the **iKeyman** utilities must be run by an administrator account.

To add the TLS certificate by using the command-line interface, 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) Issue the **iKeycmd** command to add the server's cert256.arm certificate to the Operations Center truststore.

```
ikeycmd -cert -add
-db /installation_dir/ui/Liberty/usr/servers/guiServer/gui-truststore.jks
-file /server_instance_dir/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.

server_instance_dir

The IBM Storage Protect server instance directory.

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. This directory contains the cert256.arm certificate.

If you cannot access the hub server instance directory from the **Open** window, complete the following steps:

- i) Use FTP or another file-transfer method to copy the cert256.arm files from the hub server's instance directory to the following directory on the computer where the Operations Center is installed:
 - *installation_dir*\ui\Liberty\usr\servers\guiServer
 - ii) In the **Open** window, go to the guiServer directory.
 - k) Select the cert256.arm certificate.
- Tip:** The certificate that you select must be set as the default certificate in the key database file of the hub server.
- 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.

4. Start the Operations Center web server.

5. 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. Enter the port number that is specified by either the TCPADMINPORT or SSLTCPADMINPORT 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 TLS communications between the hub server and a spoke server, see [“Securing communication between the hub server and a spoke server”](#) on page 145.

Related tasks

[“Deleting and reassigning the password for the Operations Center truststore file”](#) on page 158

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 delete the password and assign a new password.

Securing communications between the Operations Center and the hub server by using CA-signed certificates

If you use CA-signed certificates to secure the hub server, the root and intermediate CA certificate files that are sent by the certificate authority (CA) for use in the hub server must be added to the truststore file of the Operations Center.

Before you begin

Ensure that the following prerequisites are met:

- The truststore file of the Operations Center is a container for certificates that the Operations Center can access. During the installation of the Operations Center, you must create a password for the truststore file. To secure communications 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 must now re-create and configure the truststore file. For instructions, see [“Deleting and reassigning the password for the Operations Center truststore file”](#) on page 158.
- You have received the CA-signed certificates that are needed to connect to the server from the certificate authority and installed them on the server. See [Configuring the server to accept SSL connections](#).

The following figure illustrates the components for setting up a Secure Sockets Layer (SSL) connection between the hub server and the Operations Center.



About this task

To import the root and intermediate CA certificates for each IBM Storage Protect server from the hub server to the Operations Center, complete the following steps.

Tip: If you use self-signed certificates, which are the installed by default, see [“Securing communications between the Operations Center and the hub server by using self-signed certificates”](#) on page 142.

Procedure

1. Navigate to the command line of the operating system on which the Operations Center is installed.
2. From the command line, change the directory to the keystore location:
`installation_dir\ui\Liberty\usr\servers\guiServer`
 where *installation_dir* represents the directory in which the Operations Center is installed.
3. Copy the root CA certificate and intermediate CA certificate files to this location.
Tip: The certificate files were previously copied to the hub server location.
4. Stop the Operations Center web server as described in [“Starting and stopping the web server”](#) on page 160.
5. Make a backup copy of the Operations Center truststore file in case you must revert to the original version. The Operations Center truststore file is named `gui-truststore.jks`.
6. To complete the steps to receive the CA-signed certificate, use one of the following commands:
 - **ikeyman** command: See [“Receiving the signed certificate by using IBM Key Management”](#) on page 151 and go to the steps for receiving the signed certificate.
 - **ikeycmd** command: See [“Receiving the signed certificate by using ikeycmd”](#) on page 157 and go to the steps for receiving the signed certificate.
7. Start the Operations Center web server.

What to do next

To set up TLS communications between the hub server and a spoke server, follow the instructions in [“Securing communication between the hub server and a spoke server”](#) on page 145.

Related tasks

[“Receiving the signed certificate”](#) on page 151

The CA must send you the certificate file to add to the truststore file.

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, and the certificate of the hub server to the spoke server. You must also configure the Operations Center to monitor the spoke server.

About this task

The hub server receives status and alert information from the spoke server and shows this information in the Operations Center. To receive the status and alert information from the spoke server, the certificate of the spoke server must be added to the truststore file of the hub server. You must also configure the Operations Center to monitor the spoke server.

To enable other functions of the Operations Center, such as the automatic deployment of client updates, the certificate of the hub server must be added to the truststore file of the spoke server.

Procedure

1. Complete the following steps to define the certificate of the spoke server to the hub server:

- a) On the spoke server, change to the directory of the spoke server instance.
- b) Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

- c) Securely transfer the `cert256.arm` file of the spoke server to the hub server.
- d) On the hub server, change to the directory of the hub server instance.
- e) 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 -trust enable  
-label spoke_servername -file spoke_cert256.arm
```

2. Complete the following steps to define the certificate of the hub server to the spoke server:

- a) On the hub server, change to the directory of the hub server instance.
- b) Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

- c) Securely transfer the `cert256.arm` file of the hub server to the spoke server.
- d) On the spoke server, change to the directory of the spoke server instance.
- e) Define the hub server certificate to the spoke server. Issue the following command from the spoke server instance directory, where *hub_servername* is the name of the hub server, and *hub_cert256.arm* is the file name of the hub server certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii -trust enable  
-label hub_servername -file hub_cert256.arm
```

3. Restart the hub server and the spoke server.
4. Complete the following steps to define the spoke server to the hub server, and the hub server to the spoke server.
 - a) Issue the following commands on both the hub server and the spoke server:

```
SET SERVERPASSWORD server_password  
SET SERVERHLADDRESS ip_address  
SET SERVERLLADDRESS tcp_port
```

- b) On 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
```

- c) On the spoke server, issue the **DEFINE SERVER** command, according to the following example:

```
DEFINE SERVER hub_servername HLA=hub_address  
LLA=hub_SSLTCPADMINPort SERVERPA=hub_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.

5. Complete the following steps to configure the Operations Center to monitor the spoke server:

- a) On the Operations Center menu bar, click **Servers**.

The spoke server has a status of "Unmonitored." This status means that, although this server was defined to the hub server by using the **DEFINE SERVER** command, the server is not yet configured as a spoke.

- b) Click the spoke server to highlight the item, and click **Monitor Spoke**.

Configuring SSL communication between the Operations Center and web browsers

During the installation of the Operations Center, a self-signed digital certificate is generated and is then used for web browser sessions. You can optionally use a certificate that is signed by a third-party certificate authority instead of the self-signed certificate.

About this task

The Operations Center always uses the HTTPS protocol to communicate with web browsers. All communication between your browser and the Operations Center is encrypted by using version 1.2 or later of the TLS protocol.

By default, the self-signed certificate is used to create the secure connection between the browser and the Operations Center. Because the certificate is a self-signed certificate, the web browser is unable to verify the identity of the server and displays a warning. Self-signed certificates are commonly used for intranet web sites, where the danger of an intercepted connection and an impersonated server might not be considered a serious threat. You can bypass the browser's security warning and use the self-signed certificate, or you can replace the self-signed certificate with a certificate from a trusted certificate authority (CA).

To use the self-signed certificate, no further configuration is necessary.

To use a certificate that is signed by a CA, you must complete multiple steps.

Procedure

1. Create a certificate signing request.
2. Send the certificate signing request to the certificate authority for signing.
3. Add the certificate to the truststore file of the Operations Center.

Creating a certificate signing request

To get a certificate that is signed by a third party, you must create a certificate signing request (CSR) to send to the CA.

Before you begin

The truststore file of the Operations Center is a container for SSL/TLS 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 work with the truststore file, you must know the truststore password. If you do not remember this password, follow the instructions in [“Deleting and reassigning the password for the Operations Center truststore file”](#) on page 158.

Procedure

To create a CSR, complete the following steps:

1. From the command line, change the directory to the keystore location:
`installation_dir\ui\Liberty\usr\servers\guiServer`
2. Create a certificate request by using the **ikeyman** command or the **ikeycmd** command. The **ikeyman** command opens the IBM Key Management graphical user interface, and **ikeycmd** is a command-line interface.


Tip: You might have to specify the full path to the **ikeyman** or **ikeycmd** command. The commands are located in the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

`installation_dir\ui\jre\bin`


- To create a certificate request by using the **ikeman** graphical user interface, complete the following steps:
 - a. Open the IBM Key Management tool by issuing the following command:

```
ikeman
```


- b. Click **Key Database File > Open**.



In the **Open** window, click **Browse** to open the directory and select the `gui-truststore.jks` file. Click **OK**.



c. Create a certificate request. In the **Key database content** area, click **New**.



d. In the Create New Key and Certificate Request dialog box, complete the fields as required by the CA and your organization. Specify the following information:

Key Label

Specify a unique label for the certificate in the truststore file. The label name, for example, *usr-cert-name*, identifies the certificate in the truststore.

Key Size

Select a key size of at least 2048 bits.

Signature Algorithm

Select **SHA256WithRSA**.


Common Name

Specify the fully qualified domain name (FQDN) of the system on the network where the Operations Center is installed.

Remember: The FQDN for the system on your network is used in the URL for the Operations Center on your system. The URL is used by a web browser to access the Operations Center.

Enter the name of a file in which to store the certificate request

Specify a file that is named *certreq.csr* in the *guiServer* directory.



e. Close the **Open** window.

- To create a certificate request by using the **ikeycmd** command, issue the following command:

```
ikeycmd -certreq -create -db gui-truststore.jks -size 2048  
-sig_alg SHA256WithRSA -dn "CN=myhost.example.com" -file certreq.csr -label usr-cert-name  
-san_dnsname myhost.example.com,myhost  
-san_ipaddr 192.0.2.1,192.0.2.2
```

where:

-dn "CN=myhost.example.com"

Specifies the distinguished name. Input as a quoted string that contains the specification *CN=myhost.example.com*, where *myhost.example.com* specifies the FQDN of the system on the network where the Operations Center is installed.

Remember: The FQDN for the system on your network is used in the URL for the Operations Center on your system. The URL is used by a web browser to access the Operations Center.

-label *usr-cert-name*

Specifies a unique label, *usr-cert-name*, for the certificate in the truststore file.

-san_dnsname *myhost.example.com,myhost* (Optional)

Specifies the domain name server (DNS) names of the system where the Operations Center is installed. The CN and dnsname are typically the same value.

-san_ipaddr *192.0.2.1,192.0.2.2* (Optional)

Specifies the IP address of the system where the Operations Center is installed.

Sending the certificate signing request to the certificate authority

After you create the certificate request file (`certreq.csr`), you must send it to the CA for signing. Follow the instructions from the CA.

Receiving the signed certificate

The CA must send you the certificate file to add to the truststore file.

Procedure

To receive the signed certificate, complete the following steps:

1. From the command line, change the directory to the keystore location:
`installation_dir\ui\Liberty\usr\servers\guiServer`
2. Copy the files that you received from the CA to this location. These files include the CA root certificate, intermediate CA certificates (if any), and the signed certificate for the Operations Center.
3. Stop the Operations Center web server as described in [“Starting and stopping the web server”](#) on page 160.
4. Make a backup copy of the Operations Center truststore in case you must revert to the original truststore. The Operations Center truststore is named `gui-truststore.jks`.
5. To complete the steps to receive the signed certificate, use one of the following commands:
 - **ikeyman** command: Complete the steps in [“Receiving the signed certificate by using IBM Key Management”](#) on page 151.
 - **ikeycmd** command: Complete the steps in [“Receiving the signed certificate by using ikeycmd”](#) on page 157.

Receiving the signed certificate by using IBM Key Management

You can use a graphical user interface, the IBM Key Management tool, to manage the certificate keys and receive the signed certificate.

Procedure

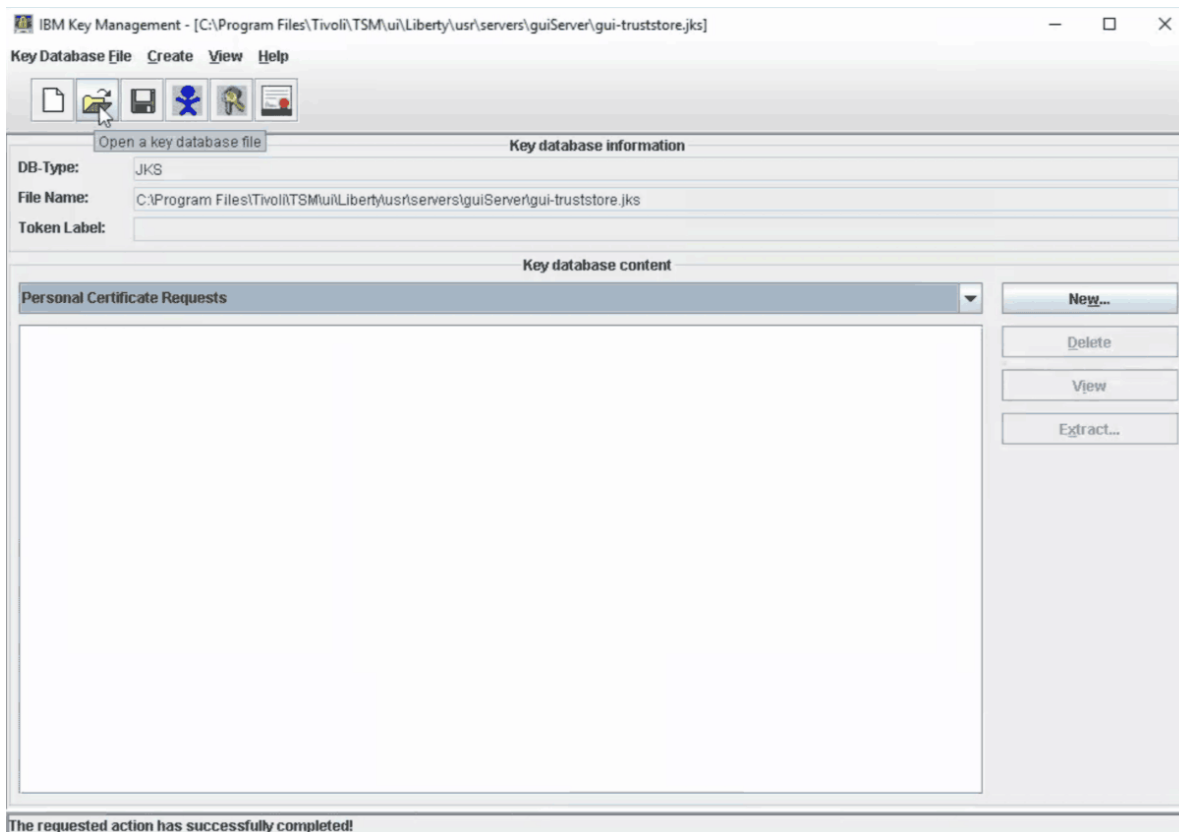
1. Verify that the Personal Signed Certificate is in the appropriate directory by using the **ikeyman** command. Complete the following steps:
 - a) Open the IBM Key Management tool by issuing the following command:

```
ikeyman
```


Tip: You might have to specify the full path to the **ikeyman** command. The commands are located in the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

```
installation_dir\ui\jre\bin
```


- b) Click **Key Database File > Open**.



In the **Open** dialog box, click **Browse** to open the directory and select the `gui-truststore.jks` file. Click **OK**.




- c) In the **Key database content** area, select **Personal Certificate Requests**, and confirm that the **usr-cert-name** label is displayed.




2. Add the CA root certificate and any intermediate certificates to the truststore file. If you received intermediate certificates from the CA, you must add each one to the truststore file before you add the CA root certificate. Complete the following steps for each intermediate certificate and the CA root certificate.

Important: The CA sends one root certificate, the signed certificate, and possibly one or more intermediate certificates. Depending on the CA, the certificate file might be one file or multiple files. If you receive the certificate file as one file, you must extract the certificates as separate files. Contact your CA if you are unsure how to extract the certificates.


- a) In the **Key database content** area, select **Signer Certificates**, and click **Add**.




- b) In the Open dialog box, specify the CA root certificate or the intermediate certificate and click **OK**.



3. Receive the signed certificate by completing the following steps:
- In the **Key database content** area, select **Personal Certificates**, and click **Receive**.




b) In the Open dialog box, specify the signed certificate and click **OK**.



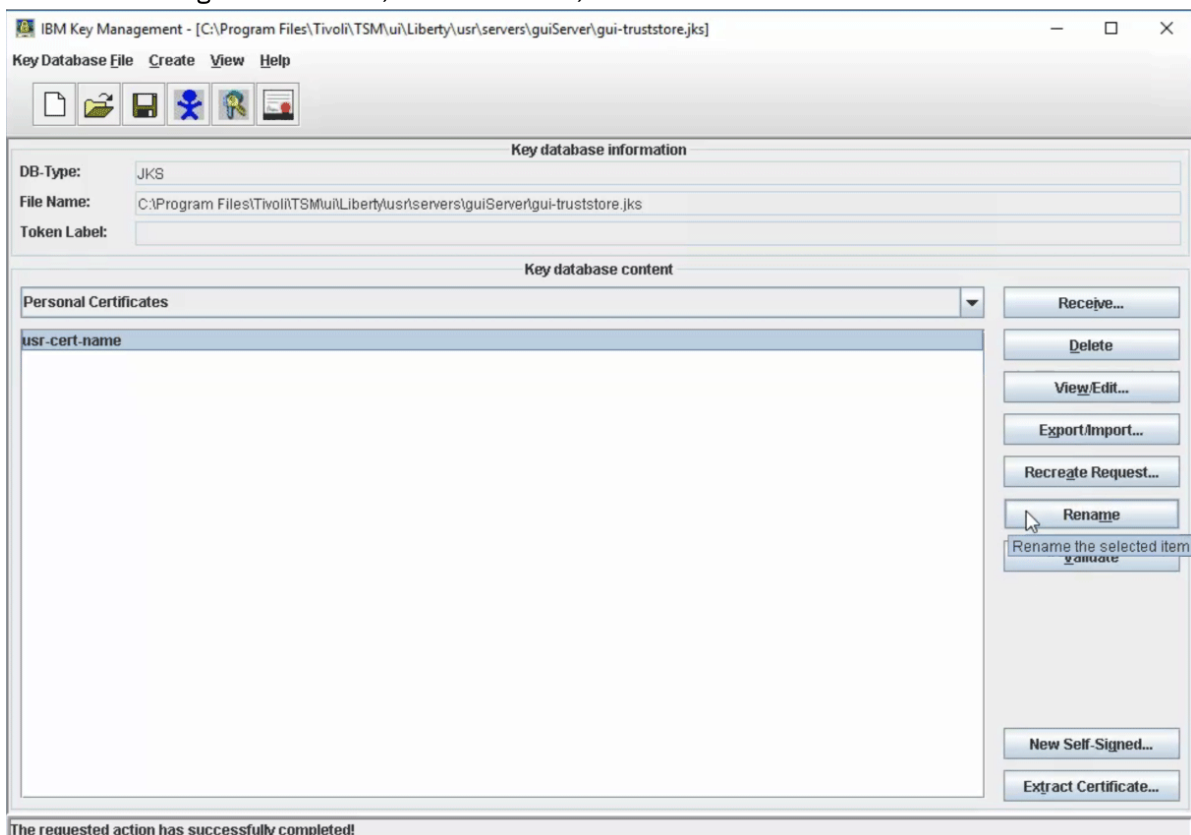
4. Delete the self-signed certificate that is currently used by the Operations Center, and replace it with the CA-signed certificate, by completing the following steps:

a) In the **Key database content** area, select **Personal Certificates**.


- b) Select the certificate that is labeled **default**, and click **Delete**. Click **Yes** in the confirmation dialog box.



- c) Select the CA-signed certificate, **usr-cert-name**, and click **Rename**.



- d) In the Rename dialog box, rename the signed certificate (usr-cert-name) to default, and click **OK**.



5. Validate the default certificate by completing the following steps:
 - a) In the **Key database content** area, select **Personal Certificates**.
 - b) Select the certificate that is labeled **default**, and click **Validate**. Click **OK** in the confirmation dialog box.
6. Start the Operations Center web server as described in [“Starting and stopping the web server”](#) on page 160.

Receiving the signed certificate by using *ikeycmd*

You can use the **ikeycmd** command, which opens a command-line, to manage certificate keys and receive signed certificates.

Procedure

1. Verify that the Personal Signed Certificate is in the appropriate directory by using the **ikeycmd** command. Complete the following steps:
 - a) Issue the following command:

```
ikeycmd -certreq -list -db gui-truststore.jks
```

Tip: You might have to specify the full path to the **ikeycmd** command. The commands are located in the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

```
installation_dir\ui\jre\bin
```

- b) A message displays the name of the Personal Signed Certificate, **usr-cert-name**, that is in the truststore file.

2. Add the CA root certificate and any intermediate certificates to the truststore file by issuing the following commands. If you received intermediate certificates from the CA, you must add them to the truststore file before you add the CA root certificate.

```
ikeycmd -cert -add -db gui-truststore.jks  
-file intermediate_certificate_file
```

```
ikeycmd -cert -add -db gui-truststore.jks  
-file root_certificate_file
```

where:

-file *certificate_file*

Specifies the name of the file that contains the certificate.

3. Receive the signed certificate by issuing the following command:

```
ikeycmd -cert -receive -db gui-truststore.jks  
-file signer_certificate_file
```

where:

-file *signer_certificate_file*

Specifies the name of the file that contains the signed certificate.

4. Delete the self-signed certificate that is currently used by the Operations Center, and replace it with the CA-signed certificate, by completing the following steps:

- a) To delete the existing self-signed certificate, issue the following command:

```
ikeycmd -cert -delete -db gui-truststore.jks -label default
```

- b) To rename the CA-signed certificate, *usr-cert-name*, to `default`, issue the following command:

```
ikeycmd -cert -rename -db gui-truststore.jks -label usr-cert-name  
-new_label default
```

where:

-label *usr-cert-name*

Identifies the CA-signed certificate by its label.

5. Validate the `default` certificate by issuing the following command:

```
ikeycmd -cert -validate -db gui-truststore.jks -label default
```

6. Start the Operations Center web server by following the instructions in [“Starting and stopping the web server”](#) on page 160.

Deleting and reassigning 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 delete the password and assign a new password.

About this task

To assign a new password, you must create a password, delete the truststore file of the Operations Center, and restart the Operations Center web server.



Attention:

If you forget the truststore password, you must get a new signed certificate from the CA. For more information, see [“Receiving the signed certificate”](#) on page 151.

Complete these steps only if you do not know the truststore password. Do not complete these steps if you know the truststore password and want to change it. To delete and reassign password, you must delete the truststore file, which deletes all certificates that are stored in the truststore file. If you know the truststore password, you can change it by using the **ikeycmd** or the **ikeyman** utility.

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 reassign the password.

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 an encrypted password as the value of the **tsm.truststore.pswd** parameter:

```
tsm.truststore.pswd={xor}MiYPPiwsKDAtoW==
```

Unencrypted password example

The following example shows an unencrypted password as the value of the **tsm.truststore.pswd** parameter:

```
tsm.truststore.pswd=J8b%^B
```

4. Replace 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.bat encode myPassword --encoding=aes
```

The following message might be shown:

```
! "java" is not recognized as an internal or external command,  
operable program or batch file.
```

If this message is shown, complete the following steps:

- i) Issue the following command, where *installation_dir* represents the directory where the Operations Center is installed:

```
set JAVA_HOME="installation_dir\ui\jre"
```

- ii) Reissue the following command to encrypt the password:

```
securityUtility.bat encode myPassword --encoding=aes
```

5. Save 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.

For information on starting the Operations Center web server, see [“Starting and stopping the web server” on page 160](#).

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.

- From the **Services** window, stop or start the Operations Center service.

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 in the range 1024 - 65535 can be specified during Operations Center installation. After installation, an administrator can configure the Operations Center to use the standard TCP/IP secure port (port 443) for HTTPS communication. If the Operations Center is configured to use port 443, then you do not need to include the secure port number when you open the Operations Center. Instead, you

can enter the following address, where *hostname* represents the name of the computer where the Operations Center is installed:

```
https://hostname/oc/
```

For more information about configuring the Operations Center to use port 443, see [“Configuring the Operations Center web server to use the standard TCP/IP secure port” on page 140.](#)

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 Storage 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

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.

Collecting diagnostic information when the Operations Center and backup-archive clients are at 8.1.13 or later

In IBM Storage Protect 8.1.13, the installation of a separate client management service package was deprecated and the feature that the client management service provided was integrated into the backup-archive client package. For installing and configuring backup-archive clients, see *Installing and configuring backup-archive clients* in IBM Documentation.

If the backup-archive client and the Operations Center are installed on your system at the 8.1.13 level, the client management service feature is no longer used by, or necessary for, the Operations Center. If a backup-archive client at the 8.1.13 level is detected, the Operations Center 8.1.13 automatically connects to the backup-archive client, not to the separate client management service.

For the Operations Center to connect to the backup-archive client, the IBM Storage Protect web user interface must be installed and configured on the workstation where the backup-archive client is installed. For instructions, see *Installing the web user interface for remote client operations* in IBM Documentation.

Note: When trying to access UNIX or Linux Backup-Archive client diagnosis data, and the client version is 8.1.13 or above, the client must use HTTPPORT 1581 (default value). If a value other than 1581 is used, the following error message is displayed. 'ANS5060E Invalid parameter passed'. To fix the error, you must use the default value (HTTPPORT 1581) in the client dsm.sys stanza and restart the dsmsvc process.

REST API endpoints

Representational State Transfer (REST) application programming interfaces (APIs) are service endpoints that support sets of HTTP operations that can create, retrieve, update, or delete access to the client management service's resources. An endpoint can send data to and receive data from the server.

If the Operations Center and the backup-archive client are at version 8.1.13 or later, client management service endpoints are in the backup-archive client web API. The Operations Center communicates directly with the API.

The REST API endpoints in the following table can be used if both the backup-archive client and the Operations Center are at 8.1.13 or later.

Table 21. REST API endpoints on 8.1.13	
REST API endpoints	Description
/ba/CMS/GetInfo	Retrieve backup-archive client configuration and log file content.
/ba/CMS/GetLogSearch	Retrieve backup-archive client log file content, filtered by a search string.
/ba/CMS/GetLogSection	Retrieve backup-archive client log file content, filtered by specifying a start line and the total number of lines.
/ba/CMS/GetLogSectionByDateTime	Retrieve backup-archive log file content, filtered by specifying a range of timestamps.

Collecting diagnostic information when the Operations Center and backup-archive client versions are earlier than 8.1.13

If the Operations Center and backup-archive client versions are earlier than 8.1.13, you must install and configure the Operations Center to access the client management service.

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.

Tip: Before you install the client management service, ensure that a successful connection was established between the backup-archive client and the server. The server truststore file that the client uses does not have the server Secure Sockets Layer (SSL) certificate until the client system has connected to the server.

You can also install the client management service on data mover nodes for IBM Storage 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

Remember: If you installed the Operations Center and backup-archive client at version 8.1.13 or later, do not install the client management service. The client management service is integrated into the backup-archive client starting with version 8.1.13.

Review [“Requirements and limitations for IBM Storage Protect client management services”](#) on page 123.

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 163](#).

- 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 Storage 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

Verify the installation.

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

Remember: If you installed the Operations Center and backup-archive client at version 8.1.13 or later, do not install the client management service. The client management service is integrated into the backup-archive client starting with version 8.1.13.

Review [“Requirements and limitations for IBM Storage Protect client management services”](#) on page 123.

Extract the installation package by following the instructions in [“Installing the client management service by using a graphical wizard”](#) on page 162.

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

Verify the installation.

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 [Configure the client management service for custom configurations](#). 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

Restriction: If you installed the Operations Center and backup-archive client at version 8.1.13 or later, you do not have to configure the Operations Center to use the client management service.

Ensure that the client management service is installed and started on the client system. Review [“Requirements and limitations for IBM Storage Protect client management services”](#) on page 123.

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 Storage 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	<code>https://server.example.com</code>
With DNS host name and non-default port	<code>https://server.example.com:1599</code>
With IP address and non-default port	<code>https://192.0.2.0:1599</code>

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 have to stop and start the service in certain situations.

Before you begin

Restriction: This procedure applies only to users who installed the Operations Center and backup-archive client at a level that is earlier than 8.1.13. If you installed the Operations Center and backup-archive client at version 8.1.13 or later, the client management service is integrated into the backup-archive client.

Procedure

- To stop, start, or restart the client management service on Linux client systems, issue the following commands:

- If the system has **systemctl** installed, issue the following commands:

- To stop the server:

```
systemctl stop cms.service
```

- To start the server:

```
systemctl start cms.service
```

- To restart the server:

```
systemctl restart cms.service
```

- To determine whether the server is running, issue the following command:

```
systemctl status cms.service
```

- If the system does not have **systemctl** installed, issue the following commands:

- To stop the server:

```
service cms.rc stop
```

- To start the server:

```
service cms.rc start
```

- To restart the server:

```
service cms.rc restart
```

- To determine whether the server is running, issue the following command:

```
service cms.rc status
```

- On Windows client systems, open the **Services** window, and stop, start, or restart the IBM Storage 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.

Before you begin

Restriction: This procedure applies only to users who installed the Operations Center and backup-archive client at a level that is earlier than 8.1.13. If you installed the Operations Center and backup-archive client at version 8.1.13 or later, the client management service is integrated into the backup-archive client.

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 Storage 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.

Before you begin

Restriction: This procedure applies only to users who installed the Operations Center, backup-archive client, and client management service at a level that is earlier than 8.1.13. If you installed the Operations

Center and backup-archive client at version 8.1.13 or later, the client management service is integrated into the backup-archive client. In this case, it is not necessary to configure the client management service for custom installations.

About this task

Before you install the client management service, ensure that a successful connection was established between the backup-archive client and the server. The server truststore file that the client uses does not have the server Secure Sockets Layer (SSL) certificate until the client system has connected to the server.

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	<code>client_install_dir/cms/bin/CmsConfig.sh</code>
Windows	<code>client_install_dir\cms\bin\CmsConfig.bat</code>

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 Storage 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 `dsmererror.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



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\dsmererror.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\dsmererror.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 Storage 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 Storage 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 Storage 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 **commethod 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 Storage 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 Storage 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 Storage 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 Storage 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 Storage 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 **CmsComflog 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 Storage 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* ➔

**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 Storage 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
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

Language	Time format for Linux client systems	Time format for Windows client systems
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
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\dsmererror.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\dsmererror.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 Storage 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 Storage 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\dsmererror.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 Storage 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 Storage 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.

Setting the log configuration of the client management services

After installing the client management service (CMS), you can set the log configuration according to your preference by editing values of the parameters in the log file. The parameters include **dateTimeFormat**, **encoding**, and **languageTag**.

Procedure

To update the parameter configuration in the log file of CMS, complete the following actions:

1. Change to the following directory where the CMS is installed: `/opt/tivoli/tsm/cms/Liberty/usr/servers/cmsServer`.
2. Open the `client-configuration.xml` file in a text editor.
3. Edit the parameters such as `dateTimeFormat`, `encoding`, and `languageTag` to your preferred values. For more information about available options, see [“CmsConfig addlog command” on page 175](#).
4. Save the file. After updating the parameters, run the `service cms.rc restart` command in the console to restart the CMS service.

5. Restart the client management service by issuing the following command: `service cms.rc restart`. For more information about starting and stopping the service, see [Starting and stopping the client management service](#).

Example

Following is an example of the `client-configuration.xml` with all the parameters.

```
<logFile>
<logPath>/opt/tivoli/tsm/client/ba/bin/dsmerror.log</logPath>
<dateTimeFormat>MM/dd/yyyy HH:mm:ss</dateTimeFormat>
<encoding>ISO-8859-1</encoding>
<languageTag>en_US</languageTag>
</logFile>
<logFile>
<logPath>/opt/tivoli/tsm/client/ba/bin/dsmsched.log</logPath>
<dateTimeFormat>MM/dd/yyyy HH:mm:ss</dateTimeFormat>
<encoding>ISO-8859-1</encoding>
<languageTag>en_US</languageTag>
</logFile>
```

In the preceding example, the parameter values are as follows:

- **dateTimeFormat** is *MM/dd/yyyy HH:mm:ss*.
- **encoding** is *ISO-8859-1*.
- **languageTag** is *en_US*.

Chapter 13. 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.
You can open IBM Installation Manager from the **Start** menu.
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:
`C:\Program Files\IBM\Installation Manager\eclipse\tools`
2. From the `tools` directory, issue the following command:
`imcl.exe -c`
3. To uninstall, enter 5.
4. Choose to uninstall from the IBM Storage 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 Storage 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:

```
C:\Program Files\IBM\Installation Manager\eclipse\tools
```

2. From the `tools` directory, issue the following command, where *response_file* represents the response file path, including the file name:

```
imcl.exe -input response_file -silent
```

The following command is an example:

```
imcl.exe -input C:\tmp\input\uninstall_response.xml -silent
```

Chapter 14. 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.

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 Storage 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 Storage Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Storage Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Documentation is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Storage Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility](http://www.ibm.com/able) (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Storage Protect family of products.

See the [IBM Storage Protect glossary](#).

Index

A

- accessibility features [189](#)
- activating
 - server
 - Windows [80](#)
- activating server options [76](#)
- active log
 - space requirements [47](#)
 - storage technology selection [38](#)
- administrative commands
 - HALT [84](#)
 - REGISTER LICENSE [84](#)
- administrator ID [125](#)
- administrator password [125](#)
- alerts
 - sending by email [137](#)
- archive failover log space
 - description [58](#)
- archive log
 - space requirements [47](#)
 - storage technology selection [38](#)
- archive log directory [71](#)

B

- backups
 - database [85](#)

C

- CA-signed certificate [144](#)
- capacity planning
 - database space requirements
 - estimates based on number of files [44](#)
 - estimates based storage pool capacity [46](#)
 - starting size [44](#)
 - recovery log space requirements
 - active and archive logs [47](#)
 - active log mirror [58](#)
- client management service
 - add log file location [175](#)
 - add node definition [171](#)
 - advanced capabilities [181](#)
 - CmsConfig addlog [175](#)
 - CmsConfig addnode [171](#)
 - CmsConfig discover [169](#)
 - CmsConfig help [181](#)
 - CmsConfig list [180](#)
 - CmsConfig remove [178](#), [179](#)
 - CmsConfig setopt [173](#)
 - CmsConfig setsys [174](#)
 - CmsConfig utility [169](#)
 - collecting diagnostic information [161](#)
 - configuring for custom client installation [168](#)
 - configuring the Operations Center [166](#)
 - installing

- client management service (*continued*)
 - installing (*continued*)
 - in silent mode [164](#)
 - Operations Center
 - view client log files [161](#)
 - remove node name [178](#), [179](#)
 - requirements and limitations [123](#)
 - REST API [181](#)
 - set client options file path [173](#)
 - set client system-options file path [174](#)
 - show configuration [180](#)
 - starting and stopping [167](#)
 - uninstalling [168](#)
 - verifying installation [165](#)
- client-configuration.xml file [165](#), [168](#), [169](#)
- cluster configuration [107](#)
- cluster environment
 - Db2 [107](#)
- clustered environment
 - applying a fix pack to a V8 server [91](#)
 - upgrading server on Windows [102](#)
 - upgrading the server [102](#)
- CmsConfig utility
 - addlog [175](#)
 - addnode [171](#)
 - client management service [169](#)
 - discover [169](#)
 - help [181](#)
 - list [180](#)
 - remove [178](#), [179](#)
 - setopt [173](#)
 - setsys [174](#)
- commands
 - administrative, SET DBRECOVERY [85](#)
 - DSMSERV FORMAT [78](#)
- commands, administrative
 - HALT [84](#)
 - REGISTER LICENSE [84](#)
- communication methods
 - Named Pipes [77](#)
 - setting [76](#)
 - TCP/IP [76](#)
- components
 - installable [vii](#)
- configuration
 - Operations Center [120](#)
- configuration wizard [73](#)
- configuring
 - communication protocols [76](#)
 - hub server [136](#)
 - Operations Center [135](#)
 - spoke server [136](#)
 - SSL [147](#)
 - TLS communication [147](#)
 - web browser communication [147](#)
- configuring the Operations Center
 - for client management service [166](#)

- configuring, manually [73, 75](#)
- configuring, server instance [73](#)
- configuring, wizard [73](#)
- Console language support [68](#)
- console mode [65](#)
- create a certificate signing request
 - third-party certificate [147](#)
- create server instance [71, 73](#)
- custom configuration
 - client management service [168](#)

D

- database
 - backups [85](#)
 - installing [78](#)
 - name [60](#)
 - storage technology selection [38](#)
- database directories [71](#)
- database manager [46](#)
- Db2 commands [109](#)
- db2icrt command [75](#)
- DEFINE DEVCLASS [85](#)
- device driver, IBM Storage Protect [vii, viii](#)
- directories
 - naming for server [60](#)
- directories, instance [71](#)
- disability [189](#)
- DISK device class
 - checklist for disk systems [35](#)
 - storage technology selection [38](#)
- disk performance
 - checklist for active log [23](#)
 - checklist for server database [21](#)
 - checklist for server recovery log [23](#)
 - checklist for storage pools on disk [35](#)
- disk space [42](#)
- disk systems
 - checklist for active log [23](#)
 - checklist for server database [21](#)
 - checklist for server recovery log [23](#)
 - classification [38](#)
 - selecting [38](#)
 - storage pools on disk [35](#)
- DSMSERV FORMAT command [78](#)
- dsmserv.v6lock [84](#)

E

- email alerts
 - suspending temporarily [139](#)
- enabling client/server communications [76](#)
- expiration
 - scheduling [33](#)

F

- FILE device class
 - checklist for disk systems [35](#)
 - storage technology selection [38](#)
- first steps [71](#)
- fix packs [89](#)
- fixes [63](#)

G

- group [71](#)

H

- HALT command [84](#)
- halting the server [84](#)
- hardware requirements
 - IBM Storage Protect [42](#)
- home directory [75](#)
- HTTPS
 - password for truststore file [126, 158](#)
- hub server
 - configuring [136](#)

I

- IBM Documentation [viii](#)
- IBM Installation Manager
 - uninstalling [115](#)
- IBM Storage Protect
 - installation [64, 65](#)
 - installation packages [63](#)
 - server changes
 - Version 8.1 [ix](#)
 - uninstalling
 - in silent mode [114](#)
 - using a graphical installation wizard [113](#)
 - using command line in console mode [113](#)
 - upgrading
 - 8.1 [95](#)
 - V7.1 to V8.1 [95](#)
- IBM Storage Protect device driver, installable package [vii, viii](#)
- IBM Storage Protect fix packs [89](#)
- IBM Storage Protect on AIX
 - upgrading
 - V8.1 [95](#)
- IBM Storage Protect support site [63](#)
- IBM Storage Protect, setting up [80](#)
- installable components [vii, viii](#)
- installation directories
 - Operations Center
 - Installation Manager [126](#)
- installation log [64, 65](#)
- Installation Manager
 - logs directory [187](#)
- installation packages
 - Operations Center [129](#)
- installation wizard [64](#)
- installing
 - client management service [162](#)
 - database [78](#)
 - device support [63](#)
 - fix packs [89](#)
 - graphical user interface
 - using [64](#)
 - minimum requirements for [42](#)
 - Operations Center [129](#)
 - recovery log [78](#)
 - server [3, 63](#)
 - using command line in console mode
 - using [65](#)

- installing (*continued*)
 - what to know about security before [3](#)
 - what to know before [3](#)
- installing the IBM Storage Protect server [67](#)
- installing the server
 - silently [67](#)
- installing Operations Center [117](#)
- instance directories [71](#)
- instance user ID [60](#)
- interim fix [89](#)
- iPad
 - monitoring the storage environment [160](#)

K

- keyboard [189](#)
- KILL command [84](#)

L

- LANGUAGE option [68](#)
- language package [69](#)
- language packages [68](#)
- language support [68](#)
- languages
 - set [68](#)
- license, IBM Storage Protect [84](#)
- licenses
 - installable package [vii](#), [viii](#)
- limitations
 - client management service [123](#)
- Local System account [81](#)
- log files
 - installation [187](#)
- login screen text
 - Operations Center [140](#)

M

- maintenance mode [83](#)
- maintenance updates [89](#)
- memory requirements [42](#)
- mobile device
 - monitoring the storage environment [160](#)
- monitoring
 - logs [86](#)
- monitoring administrator [125](#)
- multiple servers
 - upgrading
 - multiple servers [85](#)

N

- Named Pipes [77](#)
- names, best practices
 - database name [60](#)
 - directories for server [60](#)
 - instance user ID [60](#)
 - server instance [60](#)
 - server name [60](#)
- new features [ix](#)

O

- offering [42](#), [125](#)
- operating system requirements
 - Operations Center [122](#)
- Operations Center
 - administrator IDs [125](#)
 - Chrome [123](#)
 - computer requirements [120](#)
 - configuring [135](#)
 - credentials for installing [126](#)
 - Firefox [123](#)
 - hub server [120](#)
 - IE [123](#)
 - installation directory [126](#)
 - installation packages [129](#)
 - installing
 - in silent mode [130](#)
 - using a graphical wizard [129](#)
 - using command line in console mode [130](#)
 - Internet Explorer [123](#)
 - language requirements [123](#)
 - login screen text [140](#)
 - opening [136](#), [160](#)
 - operating system requirements [122](#)
 - overview [119](#)
 - password for secure communications [126](#), [158](#)
 - port number [126](#), [160](#)
 - prerequisite checks [119](#)
 - rolling back to a previous version [185](#)
 - Safari [123](#)
 - spoke server [120](#), [136](#)
 - SSL [141](#), [142](#), [144](#), [145](#)
 - standard TCP/IP secure port [140](#)
 - system requirements [119](#)
 - uninstalling
 - in silent mode [183](#)
 - using a graphical wizard [183](#)
 - using command line in console mode [183](#)
 - upgrading [117](#), [133](#)
 - URL [160](#)
 - web browser requirements [123](#)
 - web server [160](#)
- options
 - communications [76](#)
 - starting the server [80](#)
- options, client
 - SSLTCPADMINPORT [77](#)
 - SSLTCPPORT [77](#)
 - TCPADMINPORT [77](#)
 - TCPPORT [77](#)
 - TCPWINDOWSIZE [77](#)
- overview
 - Operations Center [117](#), [119](#)

P

- package [42](#), [125](#)
- package group [42](#), [125](#)
- Passport Advantage [63](#)
- password
 - encryption [131](#)
 - Operations Center [131](#)
 - Operations Center truststore file [126](#), [158](#)

- password for secure communications [126](#)
- performance
 - configuration best practices [40](#)
 - Operations Center [120](#)
- planning, capacity
 - database space requirements
 - estimates based on number of files [44](#)
 - estimates based storage pool capacity [46](#)
 - starting size [44](#)
 - recovery log space requirements
 - active log mirror [58](#)
 - recovery log space requirementsv [47](#)
- port number
 - Operations Center [126](#), [160](#)
- prerequisite checks
 - Operations Center [119](#)
- publications [viii](#)

R

- receive the signed certificate
 - IBM Key Management [151](#)
 - ikeycmd [157](#)
 - ikeyman [151](#)
 - third-party certificate [151](#), [157](#)
- recovery log
 - archive failover log space [58](#)
 - installing [78](#)
- reference, Db2 commands [109](#)
- REGISTER LICENSE command [84](#)
- Remote Execution Protocol [74](#)
- repository [42](#), [125](#)
- requirements
 - client management service [123](#)
- requirements for installation [42](#)
- resource requirements
 - Operations Center [120](#)
- reverting
 - Windows cluster [107](#)
- REXEC [74](#)
- rollback
 - Operations Center [185](#)

S

- schedule
 - daily processes [33](#)
 - server processes [33](#)
 - tuning [33](#)
- secure communications [141](#), [142](#), [144](#), [145](#)
- Secure Sockets Layer [141](#), [142](#), [144](#), [145](#)
- Secure Sockets Layer (SSL)
 - communication using [77](#)
 - retry certificate exchange [15](#)
 - Transport Layer Security (TLS) [77](#)
 - troubleshooting security updates [12](#)
 - what to know about security before you upgrade [3](#)
- send the certificate signing request
 - third-party certificate [151](#)
- server
 - naming best practices [60](#)
 - performance optimization [16](#)
 - starting

- server (*continued*)
 - starting (*continued*)
 - maintenance mode [83](#)
 - stand-alone mode [83](#)
 - starting as a service
 - configuration [81](#)
 - procedure [82](#)
 - stopping [84](#)
 - upgrading
 - to 8.1 [95](#)
 - V7.1 to V8.1 [95](#)
- server active log
 - checklist for disks [23](#)
- server AIX
 - upgrading
 - V8.1 [95](#)
- server archive log
 - checklist for disks [23](#)
- server database
 - checklist for disks [21](#)
 - directories [21](#)
 - reorganization options [79](#)
 - storage paths [21](#)
- server hardware
 - checklist for server system [17](#)
 - checklist for storage pools on disk [35](#)
 - storage technology choices [38](#)
- server instance [73](#), [75](#)
- server instance, creating [75](#)
- server instances
 - naming [60](#)
 - naming best practices [60](#)
- server license [84](#)
- server recovery log
 - checklist for disks [23](#)
- server,
 - activating [80](#)
 - setting up [80](#)
 - starting [80](#)
- server, IBM Storage Protect
 - halting [84](#)
 - options [76](#)
- services
 - starting the server as a Windows service
 - configuration [81](#)
 - procedure [82](#)
- SET DBRECOVERY [85](#)
- shared resources directory [42](#), [125](#)
- silent installation
 - IBM Storage Protect [67](#)
- software requirements
 - IBM Storage Protect [42](#)
- spoke server
 - adding [136](#)
- SSL
 - configuring [147](#)
 - password for truststore file [126](#), [158](#)
- SSL (Secure Sockets Layer)
 - communication using [77](#)
 - Transport Layer Security [77](#)
- SSLTCPADMINPORT option [77](#)
- SSLTCPPOINT option [77](#)
- stand-alone mode [83](#)
- starting

- starting (*continued*)
 - client management service [167](#)
 - server [80](#)
- startup
 - server
 - maintenance mode [83](#)
 - stand-alone mode [83](#)
- status monitoring [120](#)
- stopping
 - client management service [167](#)
 - server [84](#)
- storage pool protection
 - scheduling [33](#)
- storage pools
 - storage technology selection [38](#)
- storage technology selection [38](#)
- summary of amendments
 - Version 8.1 [ix](#)
- system requirements
 - Operations Center [119](#), [120](#), [122](#), [123](#)

T

- TCP/IP
 - setting options [76](#)
 - Version 4 [76](#)
 - Version 6 [76](#)
- TCPNODELAY option [77](#)
- TCPPORT option [77](#)
- TCPWINDOWSIZE option [77](#)
- technical changes [ix](#)
- temporary disk space [46](#)
- temporary space [46](#)
- third-party certificate
 - create a certificate signing request [147](#)
 - receive the signed certificate [151](#), [157](#)
 - send the certificate signing request [151](#)
- time
 - server upgrade [96](#)
- TLS [142](#), [144](#), [145](#)
- TLS communication
 - configuring [147](#)
- translation features [68](#)
- translations [68](#)
- Transport Layer Security (TLS) [77](#)
- Transport Layer Security protocol [142](#), [144](#), [145](#)
- truststore file
 - deleting password [158](#)
 - Operations Center [126](#)
 - reassigning password [158](#)
- tuning
 - Operations Center [120](#)

U

- Uninstall
 - IBM Installation Manager [115](#)
- uninstalling
 - client management service [168](#)
- uninstalling and reinstalling [114](#)
- updating [69](#), [133](#)
- upgrade
 - server

- upgrade (*continued*)
 - server (*continued*)
 - estimated time [96](#)
 - to 8.1 [95](#)
 - V7.1 to V8.1 [95](#)
- upgrade AIX
 - server
 - V8.1 [95](#)
- upgrading Operations Center [117](#)
- URL
 - Operations Center [160](#)
- US English [68](#)
- User Account Control [74](#)
- user ID [71](#)

V

- verifying installation
 - client management service [165](#)

W

- web server
 - starting [160](#)
 - stopping [160](#)
- Windows
 - cluster reverting [107](#)
 - system requirements [42](#)
- Windows clustered environment
 - applying a fix pack to a V8 server [91](#)
- Windows Server [74](#)
- Windows services
 - creating
 - manually [82](#)
 - starting the server
 - configuration [81](#)
 - procedure [82](#)
- wizard [71](#)
- worksheet
 - server space planning [42](#)



Product Number: 5725-W99
5725-W98
5725-X15