

IBM Spectrum Protect Plus
10.1.13

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 485.](#)

This edition applies to version 10, release 1, modification 13 of IBM Spectrum® Protect Plus (product number 5737-F11) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2017, 2022.**

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

Contents

About this publication.....	ix
Who should read this publication.....	ix
Publications	ix
Getting involved in product development.....	xi
Sponsor user program.....	xi
Beta program.....	xi
What's new in version 10.1.13.....	xiii
Chapter 1. Product overview.....	1
Deployment storyboard.....	1
Product components.....	5
Overview of the serveradmin user account.....	8
Product dashboard.....	9
Alerts.....	10
Role-based access control.....	11
Security.....	11
Incremental forever backup strategy.....	12
Replicate backup-storage data.....	13
Copying snapshots to secondary backup storage.....	14
IBM Spectrum Protect Plus on IBM Cloud.....	16
IBM Spectrum Protect Plus as a software offering on IBM Cloud.....	16
IBM Spectrum Protect Plus as a VMware service on IBM Cloud.....	17
IBM Spectrum Protect Plus on AWS.....	17
IBM Spectrum Protect Plus on Azure.....	18
Integration with IBM Cloud Pak® for Multicloud Management.....	19
Chapter 2. Installation overview.....	21
System requirements	21
Component requirements	21
Hypervisor and cloud instance requirements	21
File indexing and restore requirements.....	21
File system requirements.....	22
Db2 requirements.....	22
Microsoft Exchange Server requirements.....	22
MongoDB requirements.....	22
Microsoft 365 requirements.....	22
Oracle requirements.....	22
Microsoft SQL Server requirements.....	22
SAP HANA requirements.....	22
Post installation tasks.....	22
Assigning a static IP address.....	23
Uploading the product key.....	23
Editing firewall ports.....	24
Regenerating the Secure Sockets Layer (SSL) certificate.....	25
Verifying the Secure Sockets Layer (SSL) certificate.....	26
Chapter 3. Installing IBM Spectrum Protect Plus as a virtual appliance.....	29
Overview of virtual appliance deployment.....	29

Obtaining the IBM Spectrum Protect Plus installation package.....	29
Installing IBM Spectrum Protect Plus as a VMware virtual appliance.....	30
Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance.....	32
Chapter 4. Installing and managing vSnap servers.....	35
Installing a vSnap server.....	35
Installing a physical vSnap server.....	35
Installing a virtual vSnap server in a VMware environment.....	36
Installing a virtual vSnap server in a Hyper-V environment.....	37
Uninstalling a vSnap server.....	39
Managing vSnap servers.....	40
Registering a vSnap server.....	40
Initializing the vSnap server.....	53
Migrating onboard vSnap data to a stand-alone vSnap server.....	54
Expanding a vSnap storage pool.....	58
Changing the throughput rate.....	58
Replacing a failed vSnap server.....	59
Installing iSCSI initiator utilities.....	59
vSnap server administration reference	60
Troubleshooting vSnap servers.....	67
Chapter 5. Updating IBM Spectrum Protect Plus components.....	81
Updating IBM Spectrum Protect Plus in a virtual appliance environment.....	81
Managing updates.....	81
Updating the IBM Spectrum Protect Plus server.....	85
Updating IBM Spectrum Protect Plus in a cloud environment.....	86
Updating IBM Spectrum Protect Plus by using the user interface.....	87
Updating vSnap servers.....	88
Updating the operating system for a physical vSnap server.....	88
Updating the operating system for a virtual vSnap server.....	88
Updating a vSnap server.....	89
Additional steps for updating virtual machines in Hyper-V Replica environments.....	90
Updating VADP proxies.....	91
Applying early availability updates.....	93
Chapter 6. Getting off to a quick start.....	95
Start IBM Spectrum Protect Plus.....	96
Manage sites.....	97
Create backup policies.....	98
Create a user account for the application administrator.....	100
Back up resources.....	101
Add resources to protect.....	101
Add resources to a job definition.....	102
Back up the IBM Spectrum Protect Plus catalog.....	104
Start a backup job.....	104
Run a report.....	105
Chapter 7. Configuring the system environment.....	107
Integrating with IBM Spectrum Protect.....	107
Protecting VMware data using Open Snap Store Manager	107
Monitoring IBM Spectrum Protect Plus from the Operations Center.....	131
Managing backup storage.....	136
Managing cloud storage.....	138
Managing repository server storage.....	145
Managing sites.....	158
Adding a site.....	158
Editing a site.....	160

Deleting a site.....	162
Managing LDAP and SMTP servers.....	162
Adding an LDAP server.....	162
Adding an SMTP server.....	164
Editing settings for an LDAP or SMTP server.....	165
Deleting an LDAP or SMTP server.....	165
Managing keys and certificates for connection to IBM Spectrum Protect Plus components.....	166
Adding an access key.....	166
Deleting an access key.....	167
Adding a certificate.....	167
Deleting a certificate.....	167
Adding an SSH key.....	168
Deleting an SSH key.....	169
Managing certificates for connection to the IBM Spectrum Protect Plus user interface.....	170
Uploading a TLS certificate.....	171
Testing network connectivity.....	171
Running the Service Tool from a command line.....	171
Running the Service Tool remotely.....	172
Configuring global preferences.....	173
Configuring for virtual appliance installations.....	180
Logging on to the administrative console.....	180
Logging on to the virtual appliance.....	181
Setting the time zone.....	182
Adding virtual disks.....	182
Resetting the serveradmin password.....	186

Chapter 8. Managing SLA policies for backup operations.....189

Protection Summary.....	189
Creating an SLA policy for the IBM Spectrum Protect Plus catalog.....	192
Creating an SLA policy for the IBM Spectrum Protect Plus catalog to back up to a vSnap server..	192
Creating an SLA policy for the IBM Spectrum Protect Plus catalog to back up to cloud storage....	196
Creating an SLA policy for hypervisors.....	198
Creating an SLA policy for hypervisor backup to a vSnap server.....	198
Creating an SLA policy for VMware backup to the OSSM storage server.....	203
Creating an SLA policy for Amazon EC2 instances.....	205
Creating an SLA policy for databases and file systems.....	206
Editing an SLA policy.....	210
Deleting an SLA policy.....	211

Chapter 9. Protecting virtualized systems.....213

VMware.....	213
Adding a vCenter Server instance.....	213
Editing properties of a vCenter Server instance.....	215
vCenter Server privileges.....	216
Detecting VMware resources.....	219
Testing the connection to a vCenter Server virtual machine.....	219
Backing up VMware data.....	220
Managing VADP backup proxies.....	226
Restoring VMware data.....	232
Enabling transport encryption for VMware data.....	242
Hyper-V.....	242
Adding a Hyper-V server.....	243
Backing up Hyper-V data.....	245
Restoring Hyper-V data.....	251
Amazon EC2.....	256
Creating an AWS IAM user.....	257
Adding an Amazon EC2 account.....	258

Backing up Amazon EC2 data.....	259
Restoring Amazon EC2 data.....	261
Restoring files.....	264
Chapter 10. Protecting file systems.....	267
Windows file systems.....	267
Prerequisites for file systems.....	267
Adding a file system.....	268
Backing up file system data.....	272
Restoring file system data	277
Chapter 11. Protecting cloud management systems.....	281
Microsoft 365.....	281
Registering with Azure Active Directory	281
Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus.....	282
Detailed process logs.....	284
Backing up Microsoft 365 data.....	285
Restoring Microsoft 365 data.....	286
Chapter 12. Protecting databases.....	289
Db2.....	289
Prerequisites for Db2.....	289
Adding a Db2 application server.....	292
Backing up Db2 data.....	296
Restoring Db2 data	300
Exchange Server.....	312
Prerequisites.....	312
Privileges.....	312
Adding an Exchange application server.....	314
Backing up Exchange databases.....	316
Incremental forever backup strategy.....	320
Restoring Exchange databases.....	320
Accessing Exchange database files with instant access mode.....	349
MongoDB.....	352
Prerequisites for MongoDB.....	352
Adding a MongoDB application server.....	355
Backing up MongoDB data.....	359
Restoring MongoDB data	361
Oracle.....	377
Adding an Oracle application server.....	377
Backing up Oracle data.....	379
Restoring Oracle data.....	382
SQL Server.....	389
Adding an SQL Server application server.....	390
Backing up SQL Server data.....	392
Restoring SQL Server data.....	396
SAP HANA.....	404
Prerequisites for SAP HANA.....	404
Adding an SAP HANA application server.....	406
Backing up SAP HANA data.....	409
Restoring SAP HANA data.....	413
Chapter 13. Protecting IBM Spectrum Protect Plus.....	427
Backing up the catalog.....	427
Restoring the IBM Spectrum Protect Plus catalog.....	427
Restoring the catalog from a vSnap server.....	428
Restoring the catalog from a cloud storage system.....	428

Managing restore points.....	429
Expiring job sessions.....	429
Deleting resource metadata from the catalog.....	430
Chapter 14. Managing jobs and operations.....	431
Job types.....	431
Creating jobs and job schedules.....	432
Starting jobs on demand.....	433
Viewing jobs.....	434
Viewing backup job progress at the resource level.....	436
Viewing job logs.....	437
Viewing concurrent jobs.....	437
Pausing and resuming jobs.....	437
Editing jobs and job schedules.....	437
Canceling jobs.....	438
Deleting jobs.....	438
Rerunning partially completed backup jobs.....	439
Running an ad hoc backup job.....	439
Configuring scripts for backup and restore operations.....	440
Uploading a script.....	441
Adding a script to a server.....	441
Chapter 15. Managing reports and logs.....	443
Types of reports.....	443
Backup storage utilization reports.....	443
Protection reports.....	444
System reports.....	447
Running VM environment reports.....	449
Report actions.....	450
Running a report.....	450
Creating a custom report.....	451
Scheduling a report.....	452
Collecting and reviewing audit logs for actions.....	452
Chapter 16. Managing user access.....	455
Managing user resource groups.....	457
Creating a resource group.....	457
Editing a resource group.....	460
Deleting a resource group.....	461
Managing roles.....	461
Creating a role.....	463
Editing a role.....	465
Deleting a role.....	465
Managing user accounts.....	466
Creating a user account for an individual user.....	466
Creating a user account for an LDAP group.....	466
Editing user account credentials.....	467
Deleting a user account.....	468
Managing the superuser account.....	468
Managing identities.....	469
Adding an identity.....	469
Editing an identity.....	470
Deleting an identity.....	470
Chapter 17. Troubleshooting.....	471
Collecting log files for troubleshooting.....	471
How do I tier data to tape or cloud storage?	471

How does SAN work with IBM Spectrum Protect Plus and a vSnap server?	472
Troubleshooting failed backup operations for large Db2, MongoDB, and SAP HANA databases.....	473
Determining the minimum acceptable snapshot size.....	474
Configuring the guestapps.conf file.....	476
Chapter 18. Product messages.....	479
Message prefixes.....	479
Appendix A. Search guidelines.....	481
Appendix B. Accessibility.....	483
Notices.....	485
Glossary.....	489
Index.....	491

About this publication

This publication provides overview, planning, installation, and user instructions for IBM Spectrum Protect Plus.

Who should read this publication

This publication is intended for administrators and users who are responsible for implementing a backup and recovery solution with IBM Spectrum Protect Plus in one of the supported environments.

Publications

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

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

Getting involved in product development

You can influence the future of IBM Storage products by sharing your insights with the design and development teams. To get involved, join the sponsor user program or the beta program.

Sponsor user program

The IBM Storage sponsor user program allows you to work directly with designers and developers to influence the direction of products that you use.

IBM invites you to share your experience and expertise. By joining the program, you can help us to explore, and potentially implement, new product features that are important to you and your business.

Do you use an IBM Storage software product, such as IBM Spectrum Protect Plus?

Are you ready to share your vision?

Then sign up for the sponsor user program to participate in the product innovation process. In addition, as a sponsor user, you can preview upcoming storage releases and participate in beta programs to test new product features.

To join the sponsor user program or to obtain additional information, complete the following form:

[IBM Storage Sponsor User survey](#)

Your information will remain confidential and will be used by the IBM design and development teams only for product development purposes.

Beta program

The IBM Spectrum Protect Plus beta program gives you a first glance at upcoming product features and a chance to influence design changes. You can test new software in your environment and have a direct voice in the product development process.

The beta program attracts a broad range of participants, including customers, IBM Business Partners, and IBM employees.

The program offers the following benefits:

Gain access to early code and evaluate new product features and enhancements

You get access to the beta code before general availability of the product release to determine whether the new features and enhancements are a good fit for your organization. After the code is downloaded, you can run and validate the new software in your environment. You can then identify and resolve any concerns before the code is available, thus saving time and helping to prevent production issues later. When the code becomes available, you are ready to install it and take advantage of the new capabilities.

Interact with design and development teams

The product designers, architects, developers, and testers help to plan the beta release and support its participants. These experts can assist you with resolving any issues.

Become an IBM reference customer

After your positive beta experience, IBM invites you to participate in the reference program. The IBM marketing team helps you craft a message to let other potential beta testers know about your success in adopting and using early code.

Contact and enrollment information

You can enroll by completing the [IBM Spectrum Protect Plus Beta Program Signup Form](#).

What's new in version 10.1.13

IBM Spectrum Protect Plus 10.1.13 introduces new features and updates.

For a list of new features and updates in previous version 10 releases, see [IBM Spectrum Protect Plus updates](#).

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

Chapter 1. IBM Spectrum Protect Plus overview

IBM Spectrum Protect Plus is a data protection and availability solution for virtual environments and database applications that can be deployed in minutes and protect your environment within an hour.

IBM Spectrum Protect Plus can be implemented as a stand-alone solution or integrated with cloud storage or a repository server such as an IBM Spectrum Protect server for long-term data storage.

Deployment storyboard for IBM Spectrum Protect Plus

The *deployment storyboard* is designed to help you to successfully deploy IBM Spectrum Protect Plus in a production environment.

The storyboard lists each task in the required sequence and provides links to task instructions, videos, and guidelines in the [IBM Spectrum Protect Plus Blueprints](#) if applicable. The storyboard describes the expected outcome of tasks so that you can verify your progress as you deploy the product.

Before you start, review the system requirements for your environment. For more information, see [technote 304861](#).

If you installed IBM Spectrum Protect Plus as a virtual appliance, see the stories in [Table 1](#) and [Table 3](#).

Tip: If IBM Spectrum Protect Plus is installed as a virtual appliance, the steps in [Table 1](#) rely on the information in the [IBM Spectrum Protect Plus Blueprints](#) and on the functioning of the *Sizer tool*.

Stories that reference the vSnap server apply only if a vSnap server is the primary backup storage location. For information about the storage types that are available for workloads, see [“Managing backup storage” on page 136](#).

Story	Procedure	Expected outcome
Prepare for sizing your capacity requirements by downloading the Blueprints and the Sizer Tool spreadsheet.	Download the IBM Spectrum Protect Plus Blueprints . The download includes the Sizer Tool. For sizing guidelines, see Chapters 1-3 of the Blueprints.	You have the Sizer Tool spreadsheet and information you need to size your IBM Spectrum Protect Plus capacity requirements.

Table 1. IBM Spectrum Protect Plus installed as a virtual appliance (continued)

Story	Procedure	Expected outcome
<p>Size the capacity that is required for the primary storage in your environment.</p>	<p>Use the Sizer to size the primary storage.</p> <ol style="list-style-type: none"> 1. Open the downloaded <i>Sizer Tool</i> spreadsheet and enable macros. Save a copy of the spreadsheet to your local drive for primary storage. 2. Complete the Start Here sheet by specifying your choices for global options for the primary storage. 3. Open the VMware tab and enter data for the vCenter capacity that includes daily rate change and annual growth. 4. Open the HyperV tab and enter data for your Hyper-V capacity. 5. For each application that you are planning to use, open an application tab and enter data for your capacity needs. 6. When all the data is entered, click the Sizing Results tab to review the calculated results. 7. Set the preferred vSnap server size. To automatically specify the value for the vSnap storage pool size, click <i>Automatic</i>. 8. Enter the percentage of vSnap server reserve that you require. This reserve is the percentage of the vSnap server storage that is reserved for usage, restore operations, and for any reuse. 9. Open IBM Spectrum Protect Plus, and navigate to System Configuration > Global Preferences. Input the global preferences percentages as shown in the <i>Sizer Tool</i>. Use these percentages to set the following options: <ul style="list-style-type: none"> • Target free space error (percentage) • Target free space warning (percentage) 10. Review the results of the Sizer for your primary storage. Save the Sizer, but leave it open for inputting settings that are required for secondary storage. 	<p>The Sizer Tool spreadsheet helps you to calculate the sizing information for primary storage.</p> <p>You saved a copy of the Sizer sizing spreadsheet. If capacity requirements change, you can update the spreadsheet accordingly.</p> <p>You also have details about the required number and size of the vSnap servers and, optionally, the number of required VMware vStorage API for Data Protection proxies.</p> <p>You have details about an eight-year view of growth based on your input into the spreadsheet. You set global preferences for triggering warning and errors from the vSnap when it reaches a specified threshold based on percentage usage.</p>

Table 1. IBM Spectrum Protect Plus installed as a virtual appliance (continued)

Story	Procedure	Expected outcome
<p>Size the capacity that is required for the secondary storage in your environment.</p>	<p>Use the Sizer to size the secondary storage by following these steps. Refer to Chapter 5 of the Blueprints.</p> <ol style="list-style-type: none"> 1. Download the sizing spreadsheet from the Blueprints page and enable macros. Save a copy of the Sizer sheet to your local drive for secondary storage. 2. If there are any values, reset the <i>Sizer Tool</i> spreadsheet by clicking Click to reset. 3. Complete the Start Here sheet by specifying your choices for global options for the secondary storage. 4. Go to the Results tab of the primary storage <i>Sizer Tool</i> spreadsheet you previously saved. Copy the results that are listed in the Replication workload table and enter the values into the Optional Replication Input Workload table on the Start Here tab of the secondary storage Sizer Tool spreadsheet. 5. If you plan to protect application data, complete the application tabs. For example, you can specify options for copying data to object storage and replication policies. 6. Review the sizing results for your secondary storage. Save and close both Sizer Tool spreadsheets. 	<p>You have the sizing for the capacity for the secondary storage for your IBM Spectrum Protect Plus environment.</p> <p>You saved a copy of the Sizer for the secondary storage in your environment. If anything changes, you can alter the Sizer and make changes as required.</p> <p>You also have details about the vSnap server quantity for each year, the VADP proxy quantity, and the size of each vSnap server.</p> <p>You have details of an eight-year view of growth based on your inputs into the sizer. You set global preferences for triggering warning and errors from the vSnap when it reaches a percentage of usage.</p>
<p>Install or upgrade IBM Spectrum Protect Plus by using the ISO image for the version that you require. If you update the system environment, a new kernel is installed, and a restart is required.</p>	<p>Install IBM Spectrum Protect Plus, follow the instructions in “Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 30 or “Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance” on page 32.</p>	<p>IBM Spectrum Protect Plus is installed.</p>
<p>Install or upgrade the vSnap server by using the ISO image for the version that you require. If you are using data deduplication, the vSnap server restart can take up to 15 minutes.</p>	<p>Install the vSnap server, follow the instructions in “Installing a physical vSnap server” on page 35. If you are installing a virtual vSnap server, follow the instructions in “Installing a virtual vSnap server in a Hyper-V environment” on page 37.</p>	<p>The vSnap server is installed. To verify that the vSnap server is installed, run the vsnap show command.</p>

Story	Procedure	Expected outcome
Build the vSnap server with capacity that you derived from sizing by using the Blueprints and the Sizing Tool.	<ol style="list-style-type: none"> 1. Create volumes and map vSnap devices. 2. Map volumes to VM cluster. 3. Refer to the steps for setting up a virtual or physical vSnap server in the IBM Spectrum Protect Plus Blueprints. 	The vSnap server is built.
Add log space.	<p>Create a Linux® Multiple Device driver with three partitions to store the vSnap server storage cache, cloud cache, and log files. For the cloud cache, the capacity is set at 128 GB by default. If you plan to copy data to the cloud, you must increase the capacity. For physical vSnap servers copy data to cloud storage, you must create the /opt/vsnap-data file system with the required capacity.</p> <p>For more information about this step, see <i>Configuring a physical vSnap server using storage software provided RAID</i>, and <i>Chapter 7 Configuring Cloud Object Storage</i> in the IBM Spectrum Protect Plus Blueprints.</p>	You have set up log space for your virtual or physical vSnap servers.

Story	Procedure	Expected outcome
Complete post installation tasks.	<p>After you install IBM Spectrum Protect Plus, complete post-installation configuration tasks before you complete system management tasks.</p> <p>For more information and steps, see “Post installation tasks” on page 22.</p>	You are ready to complete system management tasks to configure your IBM Spectrum Protect Plus environment.
Register the vSnap server.	<p>Register the vSnap server. For more information and steps, see “Registering a vSnap server as a backup storage provider” on page 40.</p>	The vSnap server is registered and added to IBM Spectrum Protect Plus.
Initialize the vSnap server.	<p>After you install or upgrade IBM Spectrum Protect Plus, and added vSnap servers, initialize the vSnap servers. For information and steps, see “Completing a simple initialization” on page 53.</p>	Depending on your choice, the vSnap server is initialized with or without encryption.
Configure the vSnap server.	<p>Configure vSnap server storage options such as adding replication partners, see “Configuring backup storage options” on page 44.</p>	If you configured the data replication feature, replication partners are set up.
(Optional) Configure the vSnap server as a VADP proxy.	<p>If you are using a VADP proxy to optimize data movement to and from the vSnap server, you must register the vSnap server as a VADP proxy. For more instructions, see “Registering a VADP proxy on a vSnap server” on page 229.</p>	The vSnap server is configured as a VADP proxy.

Table 2. IBM Spectrum Protect Plus installed as a virtual appliance (continued)

Story	Procedure	Expected outcome
Set up the VMware environment that includes creating a vCenter, and registering a hypervisor.	To protect VMware data, you must first set up a vCenter Server. For instructions, see “Backing up and restoring VMware data” on page 213. Ensure that the required vCenter Server privileges are enabled. For more information about the required privileges, see “Virtual machine privileges ” on page 216.	A vCenter is set up with the required permissions so that you can start to protect VMware data.
Add users.	Add the users who will be required to use IBM Spectrum Protect Plus. For more information, see “Creating a user account for an individual user” on page 466 by using the Add User form on the page.	The users are added and granted permissions to operate IBM Spectrum Protect Plus.
Create a service level agreement (SLA) policy.	Set up an SLA policy or policies for your IBM Spectrum Protect Plus workloads. For more information about SLA policies, see Chapter 8, “Managing SLA policies for backup operations,” on page 189.	The SLA policies for your IBM Spectrum Protect Plus workloads are set up and you are ready to run backup jobs.
Update global preferences.	Administrators can edit the global preferences for all operations such as deduplication or encryption. For more information about global preferences, see “Configuring global preferences” on page 173.	If global preferences are set, they apply to the entire IBM Spectrum Protect Plus environment.

Video library

To learn more about the blueprints and other product features, see the [IBM Spectrum Protect Plus video library](#).

Product components

The IBM Spectrum Protect Plus solution is provided as a virtual appliance that includes storage and data movement components.

Sizing component requirements: Some environments might require more instances of these components to support greater workloads. For guidance about sizing, building, and integrating components in your IBM Spectrum Protect Plus environment, see the [IBM Spectrum Protect Plus Blueprints](#).

The following are the base components of IBM Spectrum Protect Plus:

IBM Spectrum Protect Plus server

This component manages the entire system. The server consists of several catalogs that track various system aspects such as restore points, configuration, permissions, and customizations. Typically, there is one IBM Spectrum Protect Plus server in a deployment, even if the deployment is spread across multiple locations.

Site

This component is an IBM Spectrum Protect Plus policy construct that is used to manage data placement in the environment. A site can be physical, such as a data center, or logical, such as a department or organization. IBM Spectrum Protect Plus components are assigned to sites to localize and optimize data paths. A deployment always has at least one site per physical location. The placement of backup data to a site is governed by service level agreement (SLA) policies.

If you are using a vSnap server as your primary backup storage location, the preferred method is to localize data movement to sites by placing vSnap servers and VADP proxies together at a single site.

vSnap server

This component is a pool of disk storage that receives data from production systems for data protection or reuse. The vSnap server consists of one or more disks and can be scaled up (by adding disks to increase capacity) or scaled out (by introducing multiple vSnap servers to improve overall performance).

The vSnap server is the required primary backup storage location for most, but not all, workload types in IBM Spectrum Protect Plus. For information about available primary backup storage by workload type, see [“Managing backup storage” on page 136](#).

In larger enterprise environments that use the vSnap server as the primary backup storage location, additional vSnap servers might be required. Each site can include one or more vSnap servers.

vSnap pool

This component is the logical organization of disks into a pool of storage space, which is used by the vSnap server component. This component is also referred to as a storage pool.

VADP proxy

This component is responsible for moving data from vSphere data stores to provide protection for VMware virtual machines and is required only for protection of VMware resources. Each site can include one or more VADP proxies.

User interfaces

IBM Spectrum Protect Plus provides the following interfaces for configuration, administrative, and monitoring tasks:

IBM Spectrum Protect Plus user interface

The IBM Spectrum Protect Plus user interface is the primary interface for configuring, administering, and monitoring data protection operations.

A key component of the interface is the dashboard, which provides summary information about the health of your environment. For more information about the dashboard, see [“Product dashboard” on page 9](#).

The menu bar in the user interface contains the following items:

Item	Description
IBM Spectrum Protect icon 	This icon opens IBM Spectrum Protect Operations Center to provide expanded data protection. This icon is active only when the URL is entered in the IBM Spectrum Protect Operations Center URL preference field on the Global Preferences page. For information about this preference, see “Configuring global preferences” on page 173 .
Alerts icon 	This icon opens the Alerts window. For more information about alerts, see “Alerts” on page 10 .
Help icon 	This icon opens the online help system.

Item	Description
User menu 	<p>This menu shows the name of the user who is logged on. The menu provides access to product information and what's new, quick start, and API documentation. You can also use this menu to complete tasks such as accessing logs and testing connections between IBM Spectrum Protect Plus and nodes.</p> <p>If you are logged on to IBM Spectrum Protect Plus as the superuser, you also use this menu to manage Transport Layer Security (TLS) certificates and the product license.</p> <p>The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role. There is only one IBM Spectrum Protect Plus superuser. For more information about the superuser, see “Managing the superuser account” on page 468.</p>

Depending on the browser window size, the navigation panel might push the main panel to the side for larger browser window sizes. For smaller browser window sizes, the navigation pane might overlap the main pane. You can click the collapse icon to collapse the navigation pane.

Restriction: The IBM Spectrum Protect Plus product does not follow International Components for Unicode (ICU) collation sorting for menus. Therefore, menus appear in code point order. In some languages, letters are sorted differently from code point order. As such, the sorted order of characters and words as they appear in menus when using these languages will appear out of expected order.

vSnap command-line interface

The vSnap command-line interface is a secondary interface for administering some data protection tasks. Run the **vsnap** command to access the command-line interface. The command can be invoked by the user ID `serveradmin` or any other operating system user who has vSnap administrator privileges.

Administrative console

The administrative console is available when IBM Spectrum Protect Plus is installed as a virtual appliance. The administrative console is used to complete administrative tasks such as updating, starting and stopping IBM Spectrum Protect Plus, resetting the credentials for the superuser account, changing the time zone for the application, and configuring network settings.

To log on to the administrative console, you can use the IBM Spectrum Protect Plus superuser account or the `serveradmin` user. The `serveradmin` user is used only to access the administrative console and the IBM Spectrum Protect Plus virtual appliance and is required in the following situations:

- To log on to the IBM Spectrum Protect Plus virtual appliance operating system when working with IBM Support.
- To log on to the administrative console to complete tasks such as resetting the credentials for the superuser account. For example, when the password for the superuser account is lost.

You cannot use the `serveradmin` user to log on to the IBM Spectrum Protect Plus.

Example VMware deployment

The following figure shows IBM Spectrum Protect Plus deployed in two active locations. Each location has inventory that requires protection. Location 1 has a vCenter server and two vSphere datacenters (and an inventory of virtual machines) and Location 2 has a single datacenter (and a smaller inventory of virtual machines).

The IBM Spectrum Protect Plus server is deployed in only one of the sites. VADP proxies and vSnap servers (with their corresponding disks) are deployed in each site to localize data movement in the context of the protected vSphere resources.

Bidirectional replication is configured to take place between the vSnap servers at the two sites.

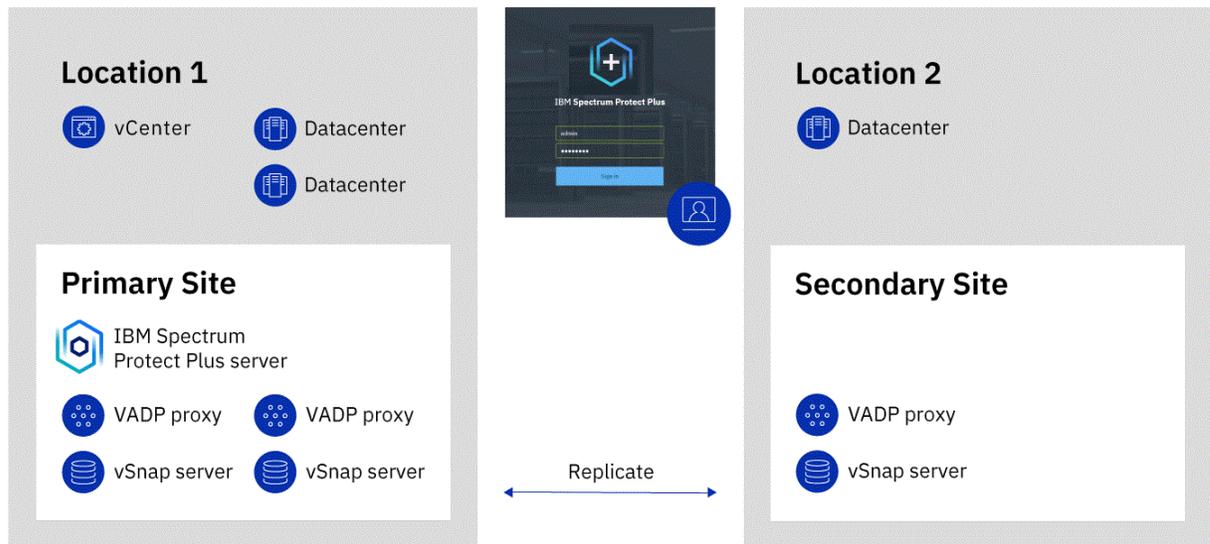


Figure 1. IBM Spectrum Protect Plus deployment across two geographical locations

Overview of the serveradmin user account

The `serveradmin` user account is a system user account that is preconfigured on IBM Spectrum Protect Plus and the vSnap server. It is used to manage both virtual appliances deployed in VMware and Microsoft Hyper-V environments.

The `serveradmin` account can be used to authenticate to the IBM Spectrum Protect Plus virtual appliance administrative console, the virtual console and using secure shell (SSH). It can also be used to access the vSnap server through the virtual console and using SSH. The initial password for the `serveradmin` user account is `sppDP758-SysXyz`. When authenticating using the `serveradmin` user account for the first time through the administrative console, the virtual console, or via SSH, you will be prompted to set a new password. For default configuration, the `serveradmin` user account password policy has these characteristics:

- The password for the user account does not expire.
- The user account is not locked after a number of failed attempts.

This may not be suitable for some environments. To harden the `serveradmin` user account password policy for IBM Spectrum Protect Plus and the vSnap server, the configuration for the account must be updated in the underlying Red Hat Enterprise Linux (RHEL) system.

Before modifying the `serveradmin` user account password properties, consider these statements:

- On vSnap servers, the operating system credentials are used for authenticating management requests from IBM Spectrum Protect Plus and for authenticating access to SMB/CIFS file shares during backup and restore operations. If you enable and configure password aging and then later change an operating system password when it expires, the change can cause interruptions to routine IBM Spectrum Protect Plus operations. Use command `vsnap user update` to change the operating system password for an account that has been used to register the vSnap server into IBM Spectrum Protect Plus. This ensures passwords used for application programming interface (API) access and SMB/CIFS access stay in sync with the operating system password.

Note: Even if you have changed the password using other means, repeat the change by running `vsnap user update` on the vSnap server.

- In IBM Spectrum Protect Plus, edit the registration of the vSnap server to update the credentials to specify the new password.
- On IBM Spectrum Protect Plus and vSnap servers, if you enable account locking for the `serveradmin` account, you may be unable to log in to the appliance if there are too many failed attempts. Depending on how you configure the account locking, the access should unlock after a certain amount of time has passed.
- You may want to log in and reset the `serveradmin` account password without waiting for the configured time to pass. This can be done through using the `root` account. By default on IBM Spectrum Protect Plus and vSnap server OVAs, the `root` account can only be accessed through the virtual console and the password for the account is unknown. The `root` account password must first be reset in order to reset the `serveradmin` account password. For more information, see [“Resetting the serveradmin password” on page 186](#).

Product dashboard

The IBM Spectrum Protect Plus dashboard summarizes the health of your virtual environment in three sections: **Jobs and Operations**, **Destinations**, and **Coverage**.

Jobs and Operations

The **Jobs and Operations** section shows a summary of job activities for a selected time period. Select the time period from the drop-down list. The following information is shown in this section:

Currently Running

The **Currently Running** section shows the total number of jobs that are running and the percentage of central processor unit (CPU) usage in the IBM Spectrum Protect Plus virtual appliance. This percentage is refreshed every 10 seconds.

To view detailed information about running jobs, click **View**.

History

The **History** section shows the total number of jobs that were completed within the selected time period. This number does not include running jobs.

This section also shows the success rate for jobs over the selected time period. The success rate is calculated by using the following formula:

$$100 \times \text{Successful Jobs} / \text{Total Jobs} = \text{Success Rate}$$

Completed jobs are shown by job status:

Successful

The number of jobs that were completed with no warnings or critical errors.

Failed

The number of jobs that failed with critical errors or that failed to be completed.

Warning

The number of jobs that were partially completed, skipped, or otherwise resulted in warnings.

To view detailed information job history information, click **View**.

Destinations

The **Destination** section shows a summary of the devices that are used for backup operations. The following information is shown in this section:

Capacity Summary

The **Capacity Summary** section shows the current usage and availability of the vSnap servers that are available to IBM Spectrum Protect Plus.

To view information about vSnap servers, click **View**.

Device Status

The **Device Status** section shows the total number of devices that are available for use.

The number of devices that are offline or otherwise unavailable is shown in the **Inactive** field.

The number of devices that are at capacity is shown in the **Full** field.

Data Reduction

The **Data Reduction** section shows data deduplication and data compression ratios.

The data deduplication ratio is the amount of data that is protected compared with the physical space that is required to store the data after duplicates are removed. This ratio represents space savings achieved in addition to the compression ratio. If deduplication is disabled, this ratio is 1.

Coverage

The **Coverage** section shows a summary of the resources that are inventoried by IBM Spectrum Protect Plus and the service level agreement (SLA) policies that are assigned to the resources. The following information is shown in this section:

Source Protection

The **Source Protection** section shows the total number of source resources, such as virtual machines and application servers, that are inventoried in the IBM Spectrum Protect Plus catalog. The number of protected and unprotected resources are shown.

This section also shows the ratio of resources that are protected in IBM Spectrum Protect Plus to the total resources, expressed as a percent.

Policies

The **Policies** section shows the total number of SLA policies with associated protection jobs.

This section also shows the three SLA policies that have the highest count assigned resources.

To view detailed information about all SLA policies, click **View**.

Alerts

The **Alerts** menu displays current and recent warnings and errors in the IBM Spectrum Protect Plus environment. The number of alerts is displayed in a red circle, indicating that alerts are available to view.

Click the **Alerts** menu to view the alerts list. Each item in the list includes a status icon, a summary of the alert, the time the associated warning or error occurred, and a link to view associated logs.

The alert list can include the following alert types:

Alert types

Job failed

Is displayed when a job fails.

Job partially succeeded

Is displayed when a job partially succeeds.

System disk space low

Is displayed when the amount of free disk space is 10% or less.

vSnap storage space low

Is displayed when the amount of free disk space is 10% or less.

System memory low

Is displayed when memory usage exceeds 95%.

System CPU usage high

Is displayed when processor usage exceeds 95%.

Hypervisor VM not found

Is displayed when the VM is not found.

Replication storage snapshot locked exception

Is displayed when the replication storage snapshot is locked. Increase replication retention or increase the replication frequency policy.

Copy storage snapshot locked exception

Is displayed when the most recently copied storage snapshot is locked. Increase copy retention or increase the copy frequency policy.

SQL log backup failure

Is displayed when log backup fails for a database.

SQL log SMO backup failure

Is displayed when there is a Server Management Object transaction log backup failure.

SQL log size too large

Is displayed when the transaction log size is larger than space available on disk.

SQL log remaining space low

Is displayed when the transaction log backup staging directory is low on disk space and displays the amount of space remaining.

Disabled deduplication on storage

Is displayed when deduplication gets disabled and displays the IP of the storage server. This will occur when the vSnap auto disable deduplication table (DDT) option is enabled and the defined size or percentage threshold is exceeded.

Role-based access control

Role-based access control defines the resources and permissions that are available to IBM Spectrum Protect Plus user accounts.

Role-based access provides users with access to only the features and resources that they require. For example, a role can allow a user to run backup and restore jobs for virtualized systems, but does not allow the user to complete administrative tasks such as creating or modifying user accounts.

To complete the tasks that are described in this documentation, the user must be assigned a role that has the required permissions. Ensure that your user account is assigned a role that has the required permissions before you start the task.

To set up and manage user access, see [Chapter 16, “Managing user access,” on page 455](#).

Creating a superuser account with the SUPERUSER role

The SUPERUSER role provides the user with access to all IBM Spectrum Protect Plus functions. The SUPERUSER role can be assigned to only one account and that account is referred to as the superuser account.

The IBM Spectrum Protect Plus administrator is prompted to create the superuser account the first time that the administrator logs on to IBM Spectrum Protect Plus. This account is automatically assigned the SUPERUSER role.

For the steps required to set the username and password for the superuser account, see [“Start IBM Spectrum Protect Plus” on page 96](#).

To manage the superuser account after it is created, see [“Managing the superuser account” on page 468](#).

Security

Security features are provided to ensure secure communication between the IBM Spectrum Protect Plus server and the IBM Spectrum Protect Plus agents for the data that you want to protect. The security method that is used depends on the agent.

Security for file system, cloud management system, and database agents

File system, cloud management system, and database agents use one of the following security verification methods. The security verification method depends on the operating system for the agent.

Linux operating systems

Provide a Secure Shell (SSH) key that matches the key of the certificate on the resource host.

Windows operating systems

Provide a Secure Sockets Layer (SSL) certificate thumbprint that matches the thumbprint of the certificate on the resource host.

For instructions about setting the security option for file systems, cloud management systems, and databases, see the topics for adding resources in [Chapter 10, “Protecting file systems,” on page 267](#), [Chapter 11, “Protecting data on cloud systems,” on page 281](#), and [Chapter 12, “Protecting databases,” on page 289](#).

Incremental forever backup strategy

IBM Spectrum Protect Plus provides a backup strategy called *incremental forever*. Rather than scheduling periodic full backup jobs, this backup solution requires only one initial full backup to the vSnap server. Afterward, an ongoing sequence of incremental backup jobs occurs.

All subsequent backup jobs back up only new or changed data from the selected resources. The backups are then reconstructed at each point in time that a backup is performed, making it possible to recover data from any single backup point.

The incremental forever backup solution provides the following advantages:

- Reduces the amount of data that goes across the network
- Reduces data growth because all incremental backups contain only the blocks that changed since the previous backup
- Reduces the duration of backup jobs

The IBM Spectrum Protect Plus incremental forever backup process creates a snapshot of selected resources. The resources that are backed up depends on the agent type.

Virtualized systems

You can back up resource data for the following virtualized systems.

VMware vCenter server instances

Back up resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters. For more information about VMware backup operations, see [“Backing up VMware data” on page 220](#).

Microsoft Hyper-V server instances

Back up resources such as VMs, virtual hard disks (VHDX), and VHDX files. For more information about Hyper-V backup operations, see [“Backing up Hyper-V data” on page 245](#).

File systems

You can back up the directories and files that are associated with physical and virtualized Windows file system servers. For more information about Windows file system backup operations, see [“Windows file systems” on page 267](#).

Cloud management systems

You can back up resources that are associated with Microsoft 365 accounts. These resources include mailboxes, calendars, contacts, and OneDrive. For more information about Microsoft 365 backup operations, see [“Backing up Microsoft 365 data” on page 285](#).

Databases

You can back up database files for the following databases. You can also backup the database logs for some databases, which enables you to select point-in-time restore options and recovery options when you create a restore job.

Db2®

For more information about Db2 backup operations, see [“Backing up Db2 data” on page 296](#).

Oracle

For more information about Oracle backup operations, see [“Backing up Oracle data” on page 379](#).

MongoDB

For more information about MongoDB operations, see [“MongoDB ” on page 352.](#)

Microsoft Exchange Server

For more information about Exchange Server backup operations, see [“Backing up Exchange databases” on page 316.](#)

Microsoft SQL Server

For more information about SQL Server backup operations, see [“Backing up SQL Server data” on page 392.](#)

SAP HANA

For more information about SAP HANA operations, see [“SAP HANA ” on page 404.](#)

Backup for Amazon EC2

Incremental forever backup operations are not used for Amazon EC2. The backup strategy for these resources is dependent on the resource type.

Amazon EC2 instances

For information about the backup configuration for Amazon EC2 instances, see [“Backing up and restoring Amazon EC2 data” on page 256.](#)

Replicate backup-storage data

When you enable replication of backup data, data from one vSnap server is asynchronously replicated to another vSnap server. For example, you can replicate backup data from a vSnap server on a primary site to a vSnap server on a secondary site.

Enabling replication of backup-storage data

Enable backup-storage data replication by taking the following actions:

1. Establish a replication partnership between vSnap servers. Replication partnerships are established in the Manage pane of a registered vSnap server. In the **Configure Storage Partners** section, select another registered vSnap server as a storage partner to serve as the target of the replication operations.

Ensure that the pool on the partner server is sufficiently large enough to hold replicated data from the primary server's pool.

2. Enable replication of backup-storage data. The replication feature is enabled by using backup policies, which are also referred to as service level agreement (SLA) policies.

These policies define parameters that are applied to backup jobs, including the frequency of backup operations and the retention policy for the backups. For more information about SLA policies, see Chapter 8, [“Managing SLA policies for backup operations,” on page 189.](#)

You can define the backup storage replication options in the **Operational Protection > Replication Policy** section of an SLA policy. Options include the frequency of the replication, the target site, and the retention of the replication.

Considerations for enabling replication of backup-storage data

Review the considerations for enabling replication of backup-storage data:

- In environments that contain more than one vSnap server, all of the vSnap servers must have a partnership established.
- If your environment includes a mixture of encrypted and unencrypted vSnap servers, select **Only use encrypted disk storage** to replicate data to encrypted vSnap servers. If this option is selected and no encrypted vSnap servers are available, the associated job will fail.
- To create one-to-many replication scenarios, where a single set of backup data is replicated to multiple vSnap servers, create multiple SLA policies for each replication site.

Copying snapshots to secondary backup storage

If your primary backup storage is a vSnap server, you can copy snapshots from the primary backup storage to secondary storage for longer-term data protection. Secondary storage is not available for container data that is backed up to cloud storage.

The following secondary backup storage targets are available for copy operations:

- IBM Cloud® Object Storage (including IBM Cloud Object Storage Systems)
- Amazon Simple Storage Service (Amazon S3)
- Microsoft Azure
- Repository servers (for the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server)

These targets support the following storage types. The storage type that you use depends on factors such as your recovery time and security goals.

Standard object storage

Standard object storage is a method of storing data in which data is stored as discrete units, or objects, in a storage pool or repository that does not use a file hierarchy but that stores all objects at the same level.

Standard object storage is an option when you copy snapshot data to an IBM Spectrum Protect server or a cloud storage system. When snapshot data is copied to standard object storage, only the most recent backup is copied. Previous backups are not transferred during cloud copy operations.

Copying snapshots to standard object storage is useful if you want relatively fast backup and recovery times and do not require the longer-term protection, cost, and security benefits that are provided by tape or cloud archive storage.

Tape or cloud archive storage

Tape storage means that data is stored on physical tape media or in a virtual tape library (VTL). Tape storage is an option when you copy snapshot data to an IBM Spectrum Protect server.

Cloud archive storage is long-term storage method that copies data to one of the following storage services: Amazon Glacier, IBM Cloud Object Storage Archive Tier, or Microsoft Azure Archive.

When you copy snapshot data to tape or to a cloud storage system, a full copy of the data is created.

Copying snapshots to tape or cloud object archive storage provides extra cost and security benefits. By storing tape volumes at a secure, offsite location that is not connected to the internet, you can help to protect your data from online threats such as malware and hackers. However, because copying to these storage types requires a full data copy, the time required to copy data increases. In addition, the recovery time can be unpredictable and the data might take longer to process before it is usable.

When you are copying data to tape from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, it is not a good idea to use the IBM Spectrum Protect tiering function. If you are archiving data to tape, you must use a cold cache storage pool. For more information about tiering, see [“How do I tier data to tape or cloud storage?”](#) on page 471. For different scenarios and more information about how to set up storage, see [“Configuration for copying or archiving data to IBM Spectrum Protect”](#) on page 145.

Adding secondary backup storage and creating backup policies

To copy snapshots to secondary storage, the following actions are required:

Action	How to
To copy snapshots to a repository server <ul style="list-style-type: none">• Set up IBM Spectrum Protect Plus as an object client in the IBM Spectrum Protect server environment.• Add the storage to IBM Spectrum Protect Plus.	See “Configuration for copying or archiving data to IBM Spectrum Protect” on page 145 and “Registering a repository server as a backup storage provider” on page 156.

Action	How to
To copy snapshots to cloud storage, add the storage to IBM Spectrum Protect Plus.	Follow the instructions for your selected storage type: <ul style="list-style-type: none"> • “Adding Amazon S3 Object Storage” on page 138 • “Adding IBM Cloud Object Storage as a backup storage provider” on page 140 • “Adding Microsoft Azure cloud storage as a backup storage provider” on page 141 • “Registering a repository server as a backup storage provider” on page 156
Create a backup policy that includes the storage.	See “Create backup policies” on page 98 .

Example deployments

The following figure shows IBM Spectrum Protect Plus deployed in two active locations. Each location has inventory that requires protection. Location 1 has a vCenter server and two vSphere datacenters (and an inventory of virtual machines) and Location 2 has a single datacenter (and a smaller inventory of virtual machines).

The IBM Spectrum Protect Plus server is deployed in only one of the sites. VADP proxies and vSnap servers (with their corresponding disks) are deployed in each site to localize data movement in the context of the protected vSphere resources.

Bi-directional replication is configured to take place between the vSnap servers at the two sites.

Snapshots are copied from the vSnap server at the secondary site to cloud storage for long-term data protection.

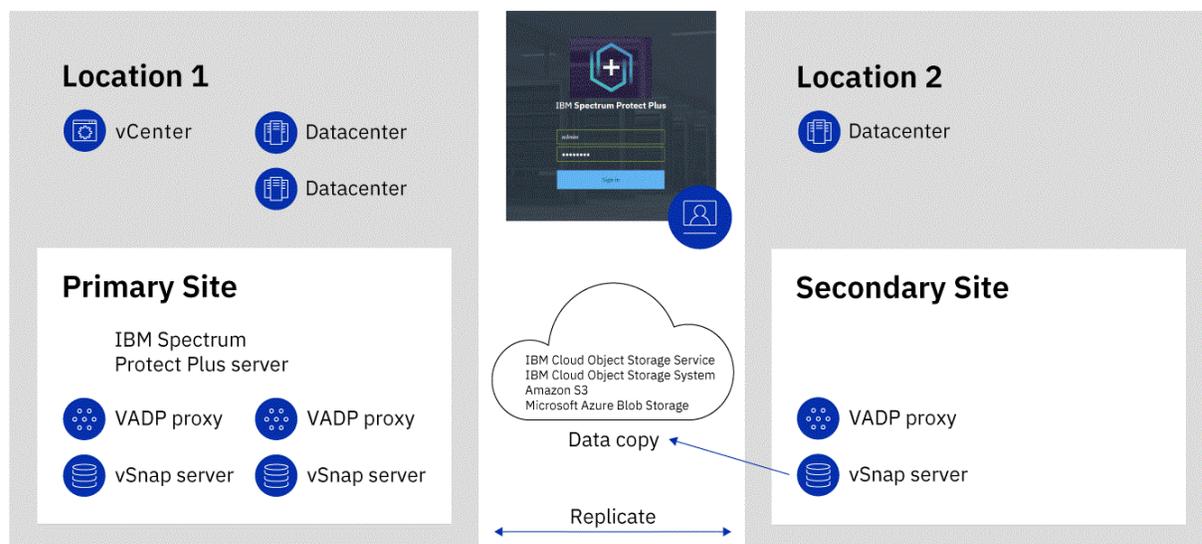


Figure 2. IBM Spectrum Protect Plus deployment across two geographical locations with copy to cloud storage

The following figure shows the same deployment as the previous figure.

However, in this deployment, snapshots are copied from the vSnap server at the secondary site to IBM Spectrum Protect for long-term data protection.

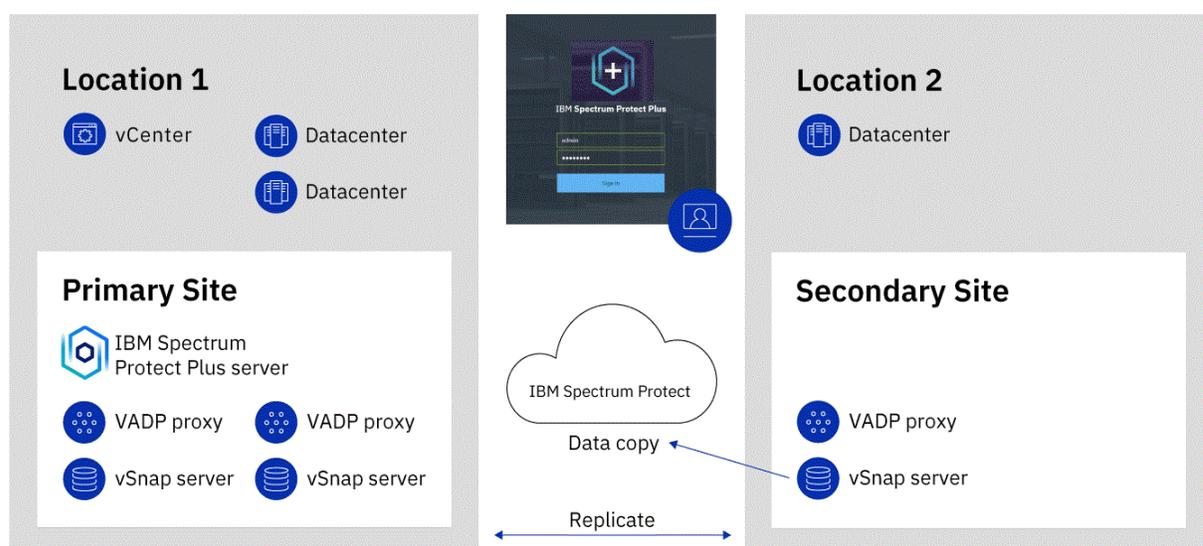


Figure 3. IBM Spectrum Protect Plus deployment across two geographical locations with copy to IBM Spectrum Protect

Related concepts

“Managing backup storage” on page 136

All IBM Spectrum Protect Plus environments must include a primary backup storage location for workload snapshots.

IBM Spectrum Protect Plus on IBM Cloud

IBM Spectrum Protect Plus is available as a software offering in the IBM Cloud catalog or as an IBM Cloud for VMware Solutions service.

IBM Spectrum Protect Plus on-premises supports both VMware and Hyper-V environments. However, IBM Spectrum Protect Plus on IBM Cloud does not support Hyper-V environments.

This documentation includes topics about features that are specific to Hyper-V. These features are not available if you are using IBM Spectrum Protect Plus on IBM Cloud.

The current version of IBM Spectrum Protect Plus and IBM Spectrum Protect Plus on IBM Cloud might not be the same.

If you want to use the current version of IBM Spectrum Protect Plus on IBM Cloud, follow the instructions in [Chapter 5, “Updating IBM Spectrum Protect Plus components,” on page 81](#) to complete an upgrade.

IBM Spectrum Protect Plus as a software offering on IBM Cloud

IBM Spectrum Protect Plus is available as a software offering in the IBM Cloud catalog.

You can deploy IBM Spectrum Protect Plus in one of the following configurations:

All-on-cloud environment

In this configuration, both the IBM Spectrum Protect Plus server and the vSnap server are deployed in IBM Cloud on an existing VMware environment. An on-premises IBM Spectrum Protect Plus server and a VMware or Microsoft Hyper-V infrastructure are not required.

This option might benefit new IBM Spectrum Protect Plus users who want to protect VMware or database applications on IBM Cloud, or existing users who are using a hybrid cloud configuration and want move to an all-on-cloud configuration.

Hybrid environment

In this configuration, only the vSnap server is deployed in IBM Cloud on an existing VMware environment. The IBM Spectrum Protect Plus server is installed and maintained on-premises. This

option might benefit existing IBM Spectrum Protect Plus users who want to continue protecting workloads that are running on premises and in the cloud environment. In addition to backup and recovery operations, you can also use a hybrid environment to replicate and reuse data between your on-premises location and IBM Cloud for more data protection. For example, you might want to use data that is protected at your on-premises site on IBM Cloud for DevOps, quality assurance, testing, and disaster recovery purposes.

Deploying IBM Spectrum Protect Plus on IBM Cloud

The IBM Spectrum Protect Plus server and vSnap server installation files are provided on separate tiles in the [IBM Cloud catalog](#).

Follow the instructions that are provided on the **Create** tab of the tile to install each server.

For videos and detailed information about IBM Cloud, see the [IBM Cloud documentation](#).

IBM Spectrum Protect Plus as a VMware service on IBM Cloud

IBM Spectrum Protect Plus is available as an IBM Cloud for VMware Solutions service.

IBM Cloud for VMware Solutions enables you to integrate or migrate your on-premises VMware workloads to the IBM Cloud by using the scalable IBM Cloud infrastructure and VMware hybrid virtualization technology.

IBM Cloud for VMware Solutions provides the following major benefits:

Global reach

Expand your hybrid cloud footprint to a maximum of 30 enterprise-class IBM Cloud datacenters around the world.

Streamlined integration

Use the streamlined process to integrate the hybrid cloud with the IBM Cloud infrastructure.

Automated deployment and configuration

Deploy an enterprise-class VMware environment with on-demand IBM Cloud Bare Metal Servers and virtual servers by using automated deployment and configuration of the VMware environment.

Simplification

Use a VMware cloud platform without identifying, procuring, deploying, and managing the underlying physical compute, storage, and network infrastructure, and software licenses.

Expansion and contraction flexibility

Expand and contract your VMware workloads according to your business requirements.

Single management console

Use a single console to deploy, access, and manage the VMware environments on IBM Cloud.

For more information

For information about how to order, install, and configure IBM Spectrum Protect Plus as a IBM Cloud for VMware Solutions service, see [Managing IBM Spectrum Protect Plus overview](#).

IBM Spectrum Protect Plus on the AWS cloud platform

IBM Spectrum Protect Plus on the Amazon Web Services (AWS) cloud platform is a data protection solution for users who want to protect databases that are running on AWS. In addition, users can protect virtual machines that are managed by VMware Cloud (VMC) on AWS while having the IBM Spectrum Protect Plus server installed on VMC and the vSnap server installed on an AWS Virtual Private Cloud (VPC).

You can deploy IBM Spectrum Protect Plus on AWS in one of the following configurations. Support for VMC on AWS is available only in a hybrid environment. For more information about support for VMC on AWS, see [IBM Spectrum Protect Plus for VMware Cloud on AWS](#).

All-on-cloud environment

In this configuration, both the IBM Spectrum Protect Plus server and the vSnap server are deployed in AWS on an existing or new VPC. An on-premises IBM Spectrum Protect Plus server and a VMware or Microsoft Hyper-V infrastructure are not required.

This option might benefit new IBM Spectrum Protect Plus users who want to protect databases on AWS and do not have IBM Spectrum Protect Plus running in an on-premises environment.

Hybrid environment

In this configuration, only the vSnap server is deployed in AWS on an existing or new VPC. The IBM Spectrum Protect Plus server is installed and maintained on premises or another location. This option might benefit existing IBM Spectrum Protect Plus users who want to continue protecting workloads that are running on premises and in the cloud environment.

In addition to backup and recovery operations, you can also use a hybrid environment to replicate and reuse data between your on-premises location and AWS for additional data protection. For example, you might want to use data that is protected at your on-premises site on AWS for DevOps, quality assurance, testing, and disaster recovery purposes.

Deploying IBM Spectrum Protect Plus to AWS

The [IBM Spectrum Protect Plus page](#) on AWS Marketplace provides the AWS CloudFormation templates that are required to deploy the IBM Spectrum Protect Plus server and vSnap server in AWS as well as pricing, usage, and support information. Follow the instructions on this page and the [IBM Spectrum Protect Plus on the AWS Cloud Deployment Guide](#) to set up your on-premises and AWS environments.

The IBM Spectrum Protect Plus on AWS deployment includes IBM Spectrum Protect Plus version 10.1.6. If you want to use the current version of IBM Spectrum Protect Plus, follow the instructions in [Chapter 5, “Updating IBM Spectrum Protect Plus components,”](#) on page 81 to complete an upgrade.

IBM Spectrum Protect Plus on the Microsoft Azure cloud platform

IBM Spectrum Protect Plus on the Microsoft Azure cloud platform is a data protection solution for users who want to protect one or more databases that are running on Azure.

IBM Spectrum Protect Plus on Azure protects the following databases and file systems that are running on Azure:

- IBM Db2
- Microsoft SQL Server
- Microsoft Exchange Server
- Oracle
- MongoDB
- Microsoft 365
- Microsoft Windows Resilient® File System (RefS) and New Technology File System (NTFS)

You can deploy IBM Spectrum Protect Plus on Azure in one of the following configurations:

All-on-cloud environment

In this configuration, both the IBM Spectrum Protect Plus server and the vSnap server are deployed in Azure on an existing or new Virtual Network (VNet).

This option might benefit new IBM Spectrum Protect Plus users who want to protect databases on Azure and do not have IBM Spectrum Protect Plus running in an on-premises environment.

Hybrid environment

In this configuration, only the vSnap server is deployed in Azure on an existing or new VNet. The IBM Spectrum Protect Plus server is installed and maintained on premises or another location. This option might benefit existing IBM Spectrum Protect Plus users who want to continue protecting workloads that are running on premises and in the cloud environment.

In addition to backup and recovery operations, you can also use a hybrid environment to replicate and reuse data between your on-premises location and Azure for additional data protection. For example, you might want to use data that is protected at your on-premises site on Azure for DevOps, quality assurance, testing, and disaster recovery purposes.

Deploying IBM Spectrum Protect Plus on Microsoft Azure

Follow the instructions in the [IBM Spectrum Protect Plus on Microsoft Azure Deployment Guide](#) to deploy IBM Spectrum Protect Plus on Azure.

Integration with IBM Cloud Pak for Multicloud Management

IBM Spectrum Protect Plus integrates with IBM Cloud Pak for Multicloud Management 2.2 or later to provide data protection for the virtual machine, and database applications in an IBM Cloud Pak for Multicloud Management environment.

IBM Cloud Pak for Multicloud Management is an open, hybrid cloud management platform that runs on the Red Hat® OpenShift® platform.

IBM Cloud Pak for Multicloud Management enables organizations to securely manage diverse applications in a hybrid cloud environment. IBM Cloud Pak for Multicloud Management provides a single control point for deploying, managing, and securing your application workloads.

For more information about IBM Cloud Pak for Multicloud Management, see the product information and demo on the [IBM Cloud Pak for Multicloud Management product page](#).

Architecture

IBM Cloud Pak for Multicloud Management is installed on an OpenShift hub cluster. IBM Spectrum Protect Plus is a partner product that is installed by using an OpenShift operator on the hub cluster.

You must have IBM Cloud Pak for Multicloud Management installed on the hub cluster prior to installing the OpenShift operator for IBM Spectrum Protect Plus.

For instructions about installing IBM Cloud Pak for Multicloud Management, go to the [online product documentation](#), select the product version, and navigate to the installation instructions.

Starting IBM Spectrum Protect Plus from IBM Cloud Pak for Multicloud Management

To start IBM Spectrum Protect Plus from IBM Cloud Pak for Multicloud Management, click **IBM Spectrum Protect Plus** on the **Administer** menu. The IBM Spectrum Protect Plus login page opens on a separate tab of the browser window.

Chapter 2. Installation overview

You can install IBM Spectrum Protect Plus as a VMware or Microsoft Hyper-V virtual appliance.

Installing as a virtual appliance

IBM Spectrum Protect Plus is installed on a VMware or Microsoft Hyper-V virtual appliance. The virtual appliance contains the application and catalogs, which manage data protection. Maintenance tasks are completed in vSphere Client or Hyper-V Manager, by using the IBM Spectrum Protect Plus command line, or in the web-based administrative console.

For more information about installing IBM Spectrum Protect Plus as a virtual appliance, see [“Overview of IBM Spectrum Protect Plus virtual appliance deployment” on page 29](#).

Installation prerequisites

Before you start the installation process, ensure that your environment meets the prerequisites that are provided in the following documents:

- [IBM Spectrum Protect Plus Blueprints](#)
- [“Deployment storyboard for IBM Spectrum Protect Plus” on page 1](#)
- [“System requirements ” on page 21](#)

System requirements

Before you install IBM Spectrum Protect Plus, review the hardware and software requirements for the product and other components that you plan to install in the storage environment.

For the system requirements, see [technote 304861](#).

To determine how to size, build, and place the components that are listed in the specifications in your IBM Spectrum Protect Plus environment, see the [IBM Spectrum Protect Plus Blueprints](#).

Component requirements

IBM Spectrum Protect Plus support for third-party platforms, applications, services, and hardware depend on the third-party vendors. When a third-party vendor product or version enters extended support, self-serve support, or end of life, IBM Spectrum Protect Plus supports the product or version at the same level as the vendor.

For the component system requirements, see [technote 6837823](#).

Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements

To help ensure that backup and restore operations for virtualized systems can run successfully, your system must meet hardware and software requirements.

For Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) system requirements, see [technote 6837825](#).

File indexing and restore requirements

To help ensure that file indexing and restore operations can run successfully, your system must meet hardware and software requirements.

For file indexing and restore system requirements, see [technote 6837847](#).

File system requirements

To help ensure that backup and restore operations for file systems can run successfully, your system must meet hardware and software requirements.

For file system requirements, see [technote 6837843](#).

Db2 requirements

To help ensure that backup and restore operations for databases can run successfully, your system must meet hardware and software requirements.

For Db2 system requirements, see [technote 6837833](#).

Microsoft Exchange Server requirements

To help ensure that backup and restore operations for Microsoft Exchange application servers can run successfully, your system must meet hardware and software requirements.

For Microsoft Exchange Server system requirements, see [technote 6837837](#).

MongoDB requirements

To help ensure that backup and restore operations for databases can run successfully, your system must meet hardware and software requirements.

For MongoDB system requirements, see [technote 6837851](#).

Microsoft 365 requirements

To help ensure that backup and restore operations for Microsoft 365 components can run successfully, your system must meet hardware and software requirements.

For Microsoft 365 system requirements, see [technote 6837853](#).

Oracle Server database backup and restore requirements

To help ensure that backup and restore operations for databases can run successfully, your system must meet hardware and software requirements.

For Oracle Server database system requirement, see [technote 6837831](#).

Microsoft SQL Server database backup and restore requirements

To help ensure that backup and restore operations for databases can run successfully, your system must meet hardware and software requirements.

For Microsoft SQL Server system requirements, see [technote 6837849](#).

SAP HANA requirements

To help ensure that backup and restore operations for databases can run successfully, your system must meet hardware and software requirements.

For SAP HANA system requirements, see [technote 6837855](#).

Post installation tasks

After you install IBM Spectrum Protect Plus, complete post-installation configuration tasks before you complete system management tasks.

Some tasks are applicable only to one type of installation: virtual appliance or as a set of OpenShift containers. Where this situation occurs, it is noted in the topic.

Assigning a static IP address

If IBM Spectrum Protect Plus is installed as a virtual appliance, a network administrator can assign a new static IP address by using the NetworkManager Text User Interface (nmtui) tool. Sudo privileges are required to run nmtui.

Procedure

To reassign a new static IP address, ensure that the IBM Spectrum Protect Plus virtual machine is powered on and complete the following steps:

1. Log on to the virtual machine console with the user ID serveradmin.
The initial password is sppDP758-SysXyz. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).
2. From a Red Hat Enterprise Linux (RHEL) command line, enter `nmtui` to open the interface.
3. From the main menu, select **Edit a connection**, and then click **OK**.
4. Select the network connection, then click **Edit**.
5. On the **Edit Connection** screen, enter an available static IP address that is not already in use.
6. Save the static IP configuration by clicking **OK**, then restart the IBM Spectrum Protect Plus appliance.

Related tasks

[“Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 30](#)

To install IBM Spectrum Protect Plus in a VMware environment, deploy an Open Virtualization Format (OVF) template. Deploying an OVF template creates a virtual appliance containing the application on a VMware host such as an ESXi server.

[“Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance” on page 32](#)

To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import the IBM Spectrum Protect Plus for Hyper-V template. Importing a template creates a virtual appliance containing the IBM Spectrum Protect Plus application on a Hyper-V virtual machine.

Uploading the product key

IBM Spectrum Protect Plus runs in a trial mode for a limited time period. A valid product key is required to use IBM Spectrum Protect Plus beyond the trial period. This product key is provided in a license file.

Before you begin

To upload the license that contains the product key, you must log on to IBM Spectrum Protect Plus as the superuser. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.

You can upload a full license or you can extend the trial by uploading a trial license. For information about available license types and how to access licenses, see [IBM Spectrum Protect Plus licenses](#) or contact your sales representative. When you obtain a license file, save the file to a computer with internet access and record the location of the file.

When a catalog backup from an IBM Spectrum Protect Plus server that is using a trial license during the evaluation period is restored to another IBM Spectrum Protect Plus server that is also using a trial license in the evaluation period, the remaining day count of the trial license of the catalog backup source server still applies. This restriction does not apply to production licenses.

Procedure

To upload a license file, complete the following steps:

1. In the IBM Spectrum Protect Plus user interface, click the user menu  in the menu bar, and then click **Upgrade license**.

The number of days until the trial license expires is shown.

2. Review the licensing information, and then click **Proceed to upload**.
3. Browse to select the license file, and then click **Upgrade license**.
4. When the license file is uploaded, close the upgrade notification window.

What to do next

After you upload the license file, complete the following action:

Action	How to
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 96 .

Editing firewall ports

Use the provided examples as a reference for opening firewall ports on remote VADP proxy servers or application servers. You must restrict port traffic to only the required network or adapters.

Use the following commands to open ports on remote VADP proxy servers or application servers.

Red Hat Enterprise Linux 7 and later, and CentOS 7 and later

Use the following command to list the open ports:

```
firewall-cmd --list-ports
```

Use the following command to list zones:

```
firewall-cmd --get-zones
```

Use the following command to list the zone that contains the Ethernet port eth0:

```
firewall-cmd --get-zone-of-interface=eth0
```

Use the following command to open port 8098 for TCP traffic. This command is not permanent.

```
firewall-cmd --add-port 8098/tcp
```

Use the following command to open port 8098 for TCP traffic after you restart the firewall rules. Use this command to make the changes persistent:

```
firewall-cmd --permanent --add-port 8098/tcp
```

To undo the change to the port, use this command:

```
firewall-cmd --remove-port 8098/tcp
```

Use the following command to open a range of ports:

```
firewall-cmd --permanent --add-port 60000-61000/tcp
```

Use the following command to reload the firewall rules with the firewall updates:

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

SUSE Linux Enterprise Server 12

Edit the SUSE Linux Enterprise Server 12 advanced security firewalls options from the **Security and Users** menu. Specify the new port range that you require and apply the changes.

Firewall configurations that use IP tables

The `iptables` utility is available on most Linux distributions to enable firewall rules and policy settings. These Linux distributions include Red Hat Enterprise Linux 6.8, Red Hat Enterprise Linux 7 and later, CentOS 7 and later, and SUSE Linux Enterprise Server 12. Before you use these commands, check which firewall zones are enabled by default. Depending upon the zone setup, the `INPUT` and `OUTPUT` terms might have to be renamed to match a zone for the required rule.

For Red Hat Enterprise Linux 7 and later, see the following example commands:

Use the following command to list the current firewall policies:

```
sudo iptables -S
```

```
sudo iptables -L
```

Use the following command to open port `8098` for inbound TCP traffic from an internal subnet `<172.31.1.0/24>`:

```
sudo iptables -A INPUT -p tcp -s 172.31.1.0/24 --dport 8098 -j ACCEPT
```

Use the following command to open port `8098` for outbound TCP traffic to internal subnet `<172.31.1.0/24>`:

```
sudo iptables -A OUTPUT -p tcp -d 172.31.1.0/24 --sport 8098 -j ACCEPT
```

Use the following command to open port `8098` for outbound TCP traffic to external subnet `<10.11.1.0/24>` and only for Ethernet port adapter `eth1`:

```
sudo iptables -A OUTPUT -o eth1 -p tcp -d 10.11.1.0/24 --sport 8098 -j ACCEPT
```

Use the following command to open port `8098` for inbound TCP traffic to a range of CES IP addresses (`10.11.1.5` through `10.11.1.11`) and only for Ethernet port adapter `eth1`:

```
sudo iptables -A INPUT -i eth1 -p tcp -m iprange --dst-range 10.11.1.5-10.11.1.11 --dport 8098 -j ACCEPT
```

Use the following command to allow an internal network, Ethernet port adapter `eth1` to communicate with an external network Ethernet port adapter `eth0`:

```
sudo iptables -A FORWARD -i eth1 -o eth0 -j ACCEPT
```

This example is for Red Hat Enterprise Linux 7 and later specifically.

Use the following command to open port `8098` for inbound traffic from subnet `10.18.0.0/24` on Ethernet port `eth1` within the public zone:

```
iptables -A IN_public_allow -i eth1 -p tcp -s 10.18.0.0/24 --dport 8098 -j ACCEPT
```

Use the following command to save firewall rule changes to persist after a firewall restart process:

```
sudo iptables-save
```

Use the following command to stop and start Uncomplicated Firewall (UFW):

```
service iptables stop service iptables start
```

Regenerating the Secure Sockets Layer (SSL) certificate

You can add or modify the IBM Spectrum Protect Plus server static network information, such as IP address or hostname after the initial deployment of IBM Spectrum Protect Plus server by regenerating

the Secure Sockets Layer (SSL) certificate thumbprint. The SSL certificate is used for performing secure operations between the IBM Spectrum Protect Plus server and host agents.

Before you begin

For Hyper-V hosted IBM Spectrum Protect Plus servers, you must configure the static network information, such as IP address or hostname after the initial deployment of IBM Spectrum Protect Plus server. You must regenerate the SSL certificate after you configure the IBM Spectrum Protect Plus server.

For VMware hosted IBM Spectrum Protect Plus servers, the IP address or hostname is specified in the settings of IBM Spectrum Protect Plus virtual appliance before the initial deployment of the server. If the IP address or hostname of the IBM Spectrum Protect Plus virtual appliance changes after the initial deployment, you must regenerate the SSL certificate.

Procedure

To regenerate SSL certificate, complete the following steps:

1. Log on to the IBM Spectrum Protect Plus console with the ID `serveradmin` by using Secure Shell (SSH) protocol, and issue the following commands:

```
export CERTFILE=/opt/ECX/virgo/configuration/ecx-beta.crt
```

```
export KEYFILE=/opt/ECX/virgo/configuration/ecx-beta.key
```

```
sudo rm $CERTFILE
```

```
sudo rm $KEYFILE
```

```
sudo /opt/ECX/tools/scripts/generate_default_cert -c $CERTFILE -k $KEYFILE
```

2. Reboot the IBM Spectrum Protect Plus appliance. After a system reboot, the certificate is ready for use.

Verifying the Secure Sockets Layer (SSL) certificate

You can verify the existing Secure Sockets Layer (SSL) certificate that is used on IBM Spectrum Protect Plus server to determine if further action is required.

About this task

To verify the SSL certificate, complete the following steps:

Procedure

1. Log on to the IBM Spectrum Protect Plus console with the ID `serveradmin` by using Secure Shell (SSH) protocol.
2. To verify the IP address, issue the following command:

```
hostname -I
```

3. To verify the certificate for IP address, issue the following command:

```
sudo openssl x509 -text -noout -in /opt/ECX/virgo/configuration/ecx-beta.crt | grep "IP Address:"
```

4. If the certificate contains the server IP address, no further action is required.

Note: Check the IP address in the certificate obtained in step (3) to ensure that it matches the IP address from the hostname obtained in step (2). For example, if the certificate contains multiple IP addresses such as IPv4, IPv6, and local addresses used by Kubernetes components in the IBM Spectrum Protect Plus server, you must look for the IPv4 address in the certificate.

5. If the IP address is not present in the certificate, you can verify the certificate for the server hostname by running the following command:

```
sudo openssl x509 -text -noout -in /opt/ECX/virgo/configuration/ecx-beta.crt | grep $HOSTNAME
```

Note: The \$HOSTNAME environment variable may contain the short hostname rather than the fully qualified domain name (FQDN).

6. You can get the machine FQDN by running the following command:

```
hostname -fqdn or hostname -A
```

7. If the certificate does not contain the value of the output of **step 6**, you must take the one of the following actions:
 - Change the server hostname to match the hostname in the certificate. The `nmtui` console mode utility allows you to change the server's hostname without rebooting the server.
 - Regenerate the SSL certificate (or regenerate if it is a CA certificate) and reboot the IBM Spectrum Protect Plus appliance. For instructions, see [“Regenerating the Secure Sockets Layer \(SSL\) certificate” on page 25](#).
 - Add an entry to the agent hosts file so that the hostname in the certificate resolves the IP address.

Chapter 3. Installing IBM Spectrum Protect Plus as a virtual appliance

Before you install IBM Spectrum Protect Plus as a virtual appliance, understand the components that are deployed, the prerequisites, and the installation procedure.

Related concepts

[“Post installation tasks” on page 22](#)

After you install IBM Spectrum Protect Plus, complete post-installation configuration tasks before you complete system management tasks.

Overview of IBM Spectrum Protect Plus virtual appliance deployment

IBM Spectrum Protect Plus can be installed as a virtual appliance. The virtual appliance contains the application and catalogs, which manage data protection.

Maintenance tasks are completed in vSphere Client or Hyper-V Manager, by using the IBM Spectrum Protect Plus command line, or in a web-based management console.

Maintenance tasks are completed by a system administrator. A system administrator is usually a senior-level user who designed or implemented the vSphere and ESXi or Hyper-V infrastructure, or a user with an understanding of IBM Spectrum Protect Plus, VMware or Hyper-V, and Linux command-line usage.

Infrastructure updates are managed by IBM update facilities. The IBM Spectrum Protect Plus user interface serves as the primary means for updating IBM Spectrum Protect Plus features and underlying infrastructure components, including the operating system and file system.

Obtaining the IBM Spectrum Protect Plus installation package

You can obtain the IBM Spectrum Protect Plus installation package from an IBM download site, such as Passport Advantage or Fix Central. These packages contain files that are required to install or update the IBM Spectrum Protect Plus components.

Before you begin

For the list of installation packages by component, and the links to the download site for the files, see [technote 6827871](#).

Procedure

Download the appropriate installation file.

A different installation file is provided for installation on VMware and Microsoft Hyper-V systems. Ensure that you download the correct file for your environment.

Important: Do not change the names of the installation or update files. The original file names are required for the installation or update process to complete without errors.

Related concepts

[“Updating IBM Spectrum Protect Plus components” on page 81](#)

You can update the IBM Spectrum Protect Plus components to get the latest features and enhancements. Software patches and updates are installed by using the IBM Spectrum Protect Plus user interface or command-line interface for these components.

Related tasks

[“Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 30](#)

To install IBM Spectrum Protect Plus in a VMware environment, deploy an Open Virtualization Format (OVF) template. Deploying an OVF template creates a virtual appliance containing the application on a VMware host such as an ESXi server.

[“Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance” on page 32](#)

To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import the IBM Spectrum Protect Plus for Hyper-V template. Importing a template creates a virtual appliance containing the IBM Spectrum Protect Plus application on a Hyper-V virtual machine.

[“Installing a vSnap server” on page 35](#)

If you are using a vSnap server as your primary backup storage location, you must have at least one vSnap server installed as part of your IBM Spectrum Protect Plus environment. The [IBM Spectrum Protect Plus Blueprints](#) will help you determine how many vSnap servers are required.

Installing IBM Spectrum Protect Plus as a VMware virtual appliance

To install IBM Spectrum Protect Plus in a VMware environment, deploy an Open Virtualization Format (OVF) template. Deploying an OVF template creates a virtual appliance containing the application on a VMware host such as an ESXi server.

Before you begin

Important: Beginning with IBM Spectrum Protect Plus 10.1.6, the hostname for the OVA deployment must not use a name that includes an underscore (_).

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in [“Component requirements ” on page 21](#) and [“Hypervisor \(Microsoft Hyper-V and VMware\) and cloud instance \(Amazon EC2\) backup and restore requirements ” on page 21](#).
- Download the virtual appliance template installation file `<part_number>.ova` from Passport Advantage® Online. For information about downloading files, see [technote 6827871](#).
- Verify the MD5 checksum of the downloaded template installation file. Ensure that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- During deployment, you will be prompted to enter network properties from the VMware user interface. You can enter a static IP address configuration, or leave all fields blank to use a DHCP configuration.
- To reassign a static IP address after deployment, you can use the NetworkManager Text User Interface (nmtui) tool. For more information, see [“Assigning a static IP address” on page 23](#).

Note the following considerations:

- You might need to configure an IP address pool that is associated with the VM network where you plan to deploy IBM Spectrum Protect Plus. Correct configuration of the IP address pool includes the setup of IP address range (if used), netmask, gateway, DNS search string, and a DNS server IP address.
- If IBM Spectrum Protect Plus will be used to protect Amazon EC2 workloads, the specified hostname must consist of only lowercase, alphanumeric characters.
- If the hostname of the IBM Spectrum Protect Plus appliance changes after deployment, either through user intervention or if a new IP address is acquired through DNS, the IBM Spectrum Protect Plus appliance must be restarted.
- A default gateway must be configured properly before deployment. Multiple DNS strings are supported, and must be separated by commas without the use of spaces.
- For later versions of vSphere, the vSphere Web Client might be required to deploy IBM Spectrum Protect Plus virtual appliances.
- IBM Spectrum Protect Plus has not been tested for IPv6 environments.

Note: Both the IBM Spectrum Protect Plus virtual appliance and the vSnap server are closed systems and anti-virus (AV) installation is not supported on virtual or physical deployments.

Procedure

To install IBM Spectrum Protect Plus as a virtual appliance, complete the following steps:

1. Deploy IBM Spectrum Protect Plus. Using either the vSphere Client (HTML5) or the vSphere Web Client (FLEX), from the **Actions** menu, click **Deploy OVF Template**.
2. Specify the location of the `<part_number>.ova` file and select it. Click **Next**.
3. Provide a meaningful name for the template, which becomes the name of your virtual machine. Identify an appropriate location to deploy the virtual machine. Click **Next**.
4. Select an appropriate destination to compute resource. Click **Next**.
5. Review the template details.

Important: If you are using the vSphere Web Client (FLEX), verify that `disk.enableUUID = true` presents in **Extra Configuration**. If that is not the case or if you are using the vSphere Client (HTML5), proceed with the installation steps and enable this option from the vSphere Web Client at a later time.

When deploying IBM Spectrum Protect Plus in a VMware vCenter 7.0 environment or later, you might receive a warning stating that the certificate is not trusted. The certificate that is supplied with IBM Spectrum Protect Plus is valid. Click **Ignore** to close the warning in the vCenter OVA UI wizard when you deploy IBM Spectrum Protect Plus.

If you prefer that your VMware vCenter fully validate the OVA certificate, download the chain and certificate authority (CA) certificates used for the signing of the OVA and install these in the vCenter. You must extract the chain and root certificates directly from DigiCert. The DigiCert Assured ID Root CA certificate and the DigiCert SHA2 Extended Validation Server CA certificate can be used to validate the OVA code signed image. Download the DigiCert Assured ID Root CA certificate at <https://cacerts.digicert.com/DigiCertAssuredIDRootCA.crt.pem> and the DigiCert SHA2 Extended Validation Server CA certificate at <https://cacerts.digicert.com/DigiCertSHA2ExtendedValidationServerCA.crt.pem>.

- a) Download each certificate in Privacy Enhanced Mail (PEM) format.
- b) Log in to the vCenter using the VMware administrator account.
- c) Navigate to Administration > Certificate Management.
- d) Locate Trusted Root Certificates and click the **Add** link to add the two downloaded certificates.

After the two certificates are added, you can import the OVA without having to click **Ignore** in the UI. Click **Next**.

6. Read and accept the End User License Agreement. Check **I accept all license agreements** for vSphere Client or click **Accept** for vSphere Web Client. Click **Next**.
7. Select the storage to which the virtual appliance is to be installed. The datastore of this storage must be configured with the destination host. The virtual appliance configuration file and the virtual disk files will be stored in it. Ensure the storage is large enough to accommodate the virtual appliance including the virtual disk files associated with it. Select a disk format of the virtual disks. Thick provisioning allows for better performance of the virtual appliance. Thin provisioning uses less disk space at the expense of performance. Click **Next**.
8. Select networks for the deployed template to use. Several available networks on the ESXi server might be available by clicking **Destination Network**. Select a destination network that allows you to define the appropriate IP address allocation for the virtual machine deployment. Click **Next**.
9. Enter the property values for the virtual appliance: Hostname, DNS, Default Gateway, Domain, Network IP Address and Network Prefix. A static IP address can be provided. If left blank, a dynamic IP address assigned by a DHCP server will be used. The network prefix must be entered using Classless Inter-Domain Routing (CIDR) notation where valid values are 1 - 24. Click **Next**.

Note: These properties can be configured using the NetworkManager Text User Interface (`nmtui`) tool. Additionally, information for the Search Domain field can be added using this command. For more information, see [Assigning a static IP address](#) on page 29.

10. Review your template settings. Click **Finish** to exit the wizard and to start deployment of the OVF template.
11. After the OVF template is deployed, power on your newly created VM. You can power on the VM from the vSphere Client.

Important: Wait several minutes for IBM Spectrum Protect Plus to initialize completely.

What to do next

Once the virtual appliance has been deployed, complete the following actions:

Action	How to
Connect to the console of the IBM Spectrum Protect Plus virtual appliance by using VMware Remote Console or SSH. Set up network configurations using the NetworkManager Text User Interface (nmtui).	See Assigning a static IP address on page 29.
Upload the product key.	See “Uploading the product key” on page 23.
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 96.

Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance

To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import the IBM Spectrum Protect Plus for Hyper-V template. Importing a template creates a virtual appliance containing the IBM Spectrum Protect Plus application on a Hyper-V virtual machine.

Before you begin

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in [“Component requirements ”](#) on page 21 and [“Hypervisor \(Microsoft Hyper-V and VMware\) and cloud instance \(Amazon EC2\) backup and restore requirements ”](#) on page 21.
- Download the installation file `<part_number>.exe` from Passport Advantage Online. For information about downloading files, see [technote 6827871](#).
- Review additional Hyper-V system requirements. See [System requirements for Hyper-V on Windows Server](#).
- Verify the MD5 checksum of the downloaded template installation file. Ensure that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- If IBM Spectrum Protect Plus will be used to protect Amazon EC2 workloads, the specified hostname must consist of only lowercase, alphanumeric characters.
- If the hostname of the IBM Spectrum Protect Plus virtual appliance changes after deployment, either through user intervention or if a new IP address is acquired through DNS, the IBM Spectrum Protect Plus virtual appliance must be restarted.
- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI Initiator Service running in their Services lists. Set startup type of this service to Automatic so that it starts running when the server starts.
- Administrative privileges may be required to complete certain steps during the installation process.

Note: Both the IBM Spectrum Protect Plus virtual appliance and vSnap server are closed systems and anti-virus (AV) installation is not supported on virtual or physical deployments.

Procedure

To install IBM Spectrum Protect Plus as a virtual appliance, complete the following steps:

1. Copy the `<part_number>.exe` file to your Hyper-V server.
2. Open the installer and complete the Setup Wizard.
3. Open Hyper-V Manager and select the required server.
4. From the **Actions** pane in Hyper-V Manager, click **Import Virtual Machine**. The Import Virtual Machine wizard opens. Click **Next**.
5. In the **Locate Folder** step, click **Browse** and navigate to the folder that was designated during the installation. Select the folder with **SPP-{release}** in it. Click **Next**.
6. In the **Select Virtual Machine** step, ensure the virtual machine **SPP-{release}** is selected and then click **Next**. The **Choose Import Type** dialog opens.
7. In the **Choose Import Type** step, select **Register the virtual machine in-place (use the existing unique ID)**. Click **Next**.

Important: Do not import multiple IBM Spectrum Protect Plus virtual appliances on a single Hyper-V server.

8. In the **Connect Network** step, set Connection to the virtual switch to use. Click **Next**.
9. In the **Summary** step, review the Description. Click **Finish** to close the Import Virtual Machine wizard.
10. In Hyper-V Manager, locate the new virtual machine named **SPP-{release}**. Right-click this virtual machine and click **Settings**.
11. The Settings dialog for this virtual machine will open. In the navigation panel, click **Hardware > IDE Controller 0 > Hard Drive**.
12. In the Media section, ensure that the correct virtual hard disk is selected. Note the file name of the original virtual disk. Click **Edit**.
13. The Edit Virtual Hard Disk Wizard will open. Go to the **Choose Action** step.
14. In the **Choose Action** step, click **Convert** and then click **Next**.
15. In the **Choose Disk Format** step, ensure that **VHDX** is selected. Click **Next**.
16. For the **Choose Disk Type** step, click **Fixed Size**. Click **Next**.
17. For the **Configure Disk** step, locate the folder to store the virtual disk file of the IBM Spectrum Protect Plus virtual appliance. Reuse the same file name that was noted in Step 12. If the same installation directory from Step 12 is reused, use a different name. Click **Next**.

Important: Ensure that the disk drive on which the folder resides has enough disk space available to accommodate the fixed-size virtual disk file.

18. In the **Summary** step, review the Description. Click **Finish** to close the Edit Virtual Hard Disk wizard and to initiate the conversion of the virtual disk. Once the process completes, the original virtual hard disk file may be deleted.
19. In the Settings dialog for the virtual machine, click **Browse**. Open the newly created virtual hard disk (VHDX) file that was created in the previous step.
20. Repeat steps 12 through 19 for each hard drive under **Hardware > SCSI Controller**. Click **OK** to close the Settings dialog.
21. In the Hyper-V Manager, right-click the virtual machine and click **Start**.
22. Use Hyper-V Manager to identify the IP address of the new virtual machine if the address is automatically assigned. To assign a static IP to the virtual machine, use the NetworkManager Text User Interface (nmtui) tool.

For more information, see [“Assigning a static IP address”](#) on page 23.

Important: IBM Spectrum Protect Plus or vSnap virtual machines that are deployed using Hyper-V failover clustering should be configured with a static media access control (MAC) address for each virtual network adapter. If a dynamic MAC address is used, the Linux networking configuration may be lost after failover because a new MAC address is assigned to the virtual network adapter. The MAC address may be configured by editing the settings of the virtual machine in the Hyper-V Manager or

Failover Cluster Manage. Ensuring that each virtual network adapter is assigned a static MAC address will prevent the loss of the network configuration.

23. You must regenerate the SSL certificate after you configure the IBM Spectrum Protect Plus server. To regenerate the SSL certificate, complete the following steps:

a. Log on to the IBM Spectrum Protect Plus console with the ID `serveradmin` by using Secure Shell (SSH) protocol, and issue the following commands:

i) `export CERTFILE=/opt/ECX/virgo/configuration/ecx-beta.crt`

ii) `export KEYFILE=/opt/ECX/virgo/configuration/ecx-beta.key`

iii) `sudo rm $CERTFILE`

iv) `sudo rm $KEYFILE`

v) `sudo /opt/ECX/tools/scripts/generate_default_cert -c $CERTFILE -k $KEYFILE`

b. Reboot the IBM Spectrum Protect Plus appliance. After a system reboot, the certificate is ready for use.

What to do next

After you install the virtual appliance, complete the following actions:

Action	How to
Restart the virtual appliance.	Refer to the documentation for the virtual appliance.
Upload the product key.	See “Uploading the product key” on page 23.
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 96.

Chapter 4. Installing and managing vSnap servers

The vSnap server is the required primary backup storage location for most, but not all, workload types in IBM Spectrum Protect Plus.

For information about available primary backup storage by workload type, see [“Managing backup storage” on page 136](#).

In larger enterprise environments that use the vSnap server as the primary backup storage location, additional vSnap servers might be required. For guidance about sizing, building, and placing vSnap servers and other components in your IBM Spectrum Protect Plus environment, see the [IBM Spectrum Protect Plus Blueprints](#).

Additional vSnap servers can be installed on either virtual or physical appliances any time after the IBM Spectrum Protect Plus virtual appliance is deployed. After deployment, some registration and configuration steps are required for these stand-alone vSnap servers.

The process for setting up a stand-alone vSnap server is as follows:

1. Install the vSnap server.
2. Add the vSnap server as Disk Storage in IBM Spectrum Protect Plus.
3. Initialize the system and create a storage pool.

Installing a vSnap server

If you are using a vSnap server as your primary backup storage location, you must have at least one vSnap server installed as part of your IBM Spectrum Protect Plus environment. The [IBM Spectrum Protect Plus Blueprints](#) will help you determine how many vSnap servers are required.

Before you begin

Complete the following steps:

1. Review the vSnap system requirements in [“Component requirements ” on page 21](#).
2. Download the installation package. Different installation files are provided for installation on physical or virtual machines. Ensure that you download the correct files for your environment. For more information about downloading files and other useful information, see [technote 567387](#).

Note: Both the IBM Spectrum Protect Plus virtual appliance and the vSnap server are closed systems and anti-virus (AV) installation is not supported on virtual or physical deployments.

Important: IBM Spectrum Protect Plus components, including vSnap, should not be installed on the same machine, physical or virtual, as IBM Spectrum Protect Server.

Installing a physical vSnap server

A Linux operating system that supports physical vSnap installations is required to install a vSnap server on a physical machine.

Procedure

1. Install a Linux operating system that supports physical vSnap installations.

See vSnap server physical installation information in [technote 6837823](#) for supported operating systems.

The minimum installation configuration is sufficient, but you can also install additional packages including a graphical user interface (GUI). The root partition must have at least 8 GB of free space after installation.

2. Edit the `/etc/selinux/config` file to change the SELinux mode to Permissive:

```
SELINUX=permissive
```

3. Issue the `setenforce 0` to apply the setting immediately without requiring a restart:

```
$ setenforce 0
```

4. Download the vSnap server installation file `<part_number>.run` from Passport Advantage Online. For information about downloading files, see [technote 6827871](#).

5. Make the file executable and then run the executable.

```
$ chmod +x <part_number>.run
```

6. Run the executable. The vSnap packages are installed, plus all of required components.

```
$ ./<part_number>.run
```

Alternatively, non-interactive installations or updates of vSnap may be initiated using the `noprompt` option. When this option is used, the vSnap installer will skip prompting for responses and assume an answer of "yes" to the following prompts:

- License agreement
- Kernel installation or update
- Reboot at the end of the installation or update if necessary

To use the `noprompt` option, issue the following command. Observe the deliberate space both before and after the double dashes:

```
$ sudo ./<part_number>.run -- noprompt
```

What to do next

After you install the vSnap server, complete the following action:

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See “Managing vSnap servers” on page 40 .

Installing a virtual vSnap server in a VMware environment

To install a virtual vSnap server in a VMware environment, deploy an Open Virtualization Format (OVF) template. This creates a machine that contains the vSnap server.

Before you begin

For easier network administration, use a static IP address for the virtual machine. Assign the address by using the NetworkManager Text User Interface (nmtui) tool.

For instructions, see [“Assigning a static IP address” on page 23](#), Work with your network administrator when configuring network properties.

Procedure

1. Download the vSnap server template file `<part_number>.ova` from Passport Advantage Online. For information about downloading files, see [technote 6827871](#).
2. Deploy the vSnap server. Using the vSphere Client (HTML5) or the vSphere Web Client (FLEX), click the **Actions** menu and then click **Deploy OVF Template**.
3. Specify the location of the `<part_number>.ova` file and select it. Click **Next**.

4. Provide a meaningful name for the template, which becomes the name of your virtual machine. Identify an appropriate location to deploy the virtual machine. Click **Next**.
5. Select an appropriate destination to compute resource. Click **Next**.
6. Review the template details. Click **Next**.
7. Read and accept the End User License Agreement. Check **I accept all license agreements** for vSphere Client or click **Accept** for vSphere Web Client. Click **Next**.
8. Select the storage to which the virtual appliance is to be installed. The datastore of this storage must be configured with the destination host. The virtual appliance configuration file and the virtual disk files will be stored in it. Ensure the storage is large enough to accommodate the virtual appliance including the virtual disk files associated with it. Select a disk format of the virtual disks. Thick provisioning allows for better performance of the virtual appliance. Thin provisioning uses less disk space at the expense of performance. Click **Next**.
9. Select networks for the deployed template to use. Several available networks on the ESX server may be available by clicking Destination Networks. Select a destination network that allows you to define the appropriate IP address allocation for the virtual machine deployment. Click **Next**.
10. Enter network properties for the virtual machine default gateway, DNS, search domain, IP address, network prefix, and machine host name. If you are using a Dynamic Host Configuration Protocol (DHCP) configuration, leave all fields blank.

Restriction: A default gateway must be properly configured before deployment of the OVF template. Multiple DNS strings are supported, and must be separated by commas without the use of spaces. The network prefix should be specified by a network administrator. The network prefix must be entered using CIDR notation; valid values are 1 - 24.

11. Click **Next**.
12. Review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template. Deployment might take significant time.
13. After the OVF template is deployed, power on your newly created virtual machine. You can power on the VM from the vSphere Client.

Important: It is important to keep the VM powered on.

14. Record the IP address of the newly created VM.

The IP address is required to access and register the vSnap server. Find the IP address in vSphere Client by clicking the VM and reviewing the **Summary** tab.

What to do next

After you install the vSnap server, complete the following action:

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See “Managing vSnap servers” on page 40 .
For easier network administration, assign a static IP address for the virtual machine. Use the NetworkManager Text User Interface (nmtui) tool to assign the IP address.	For instructions, see “Assigning a static IP address” on page 23 . Work with your network administrator when configuring network properties.

Installing a virtual vSnap server in a Hyper-V environment

To install a vSnap server in a Hyper-V environment, import a Hyper-V template. This creates a virtual appliance containing the vSnap server on a Hyper-V virtual machine.

Before you begin

All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI initiator service running in their Services list. Set the service to Automatic so that it is available when the machine is restarted.

Procedure

1. Download the vSnap installation file *<part_number>.exe* from Passport Advantage Online. For information about downloading files, see [technote 6827871](#).
2. Copy the installation file to your Hyper-V server.
3. Start the installer and complete the installation steps.
4. Open Hyper-V Manager and select the required server.
For Hyper-V system requirements, see [System requirements for Hyper-V on Windows Server](#).
5. From the **Actions** menu in Hyper-V Manager, click **Import Virtual Machine**, and then click **Next**. The **Locate Folder** dialog opens.
6. Browse to the location of the Virtual Machines folder within the unzipped vSnap folder. Click **Next**. The **Select Virtual Machine** dialog opens.
7. Select vSnap, and then click **Next**. The **Choose Import Type** dialog opens.
8. Choose the following import type: **Register the virtual machine in place**. Click **Next**.
9. If the Connect Network dialog opens, specify the virtual switch to use, and then click **Next**. The Completing Import dialog opens.
10. Review the description, and then click **Finish** to complete the import process and close the **Import Virtual Machine** wizard. The virtual machine is imported.
11. Right-click the newly deployed VM, and then click **Settings**.
12. Under the section named IDE Controller 0, select **Hard Drive**.
13. Click **Edit**, and then click **Next**.
14. In the **Choose Action** screen, choose **Convert** then click **Next**.
15. For the Disk Format, select **VHDX**.
16. For the Disk Type, select **Fixed Size**.
17. For the Configure Disk option, give the disk a new name and optionally, a new location.
18. Review the description, and then click **Finish** to complete the conversion.
19. Click **Browse**, and then locate and select the newly created VHDX.
20. Repeat steps 12 through 18 for each disk under the SCSI Controller section.
21. Power on the VM from **Hyper-V Manager**. If prompted, select the option where the kernel starts in rescue mode.
22. Use Hyper-V Manager to identify the IP address of the new virtual machine if automatically assigned. To assign a static IP to the virtual machine using NetworkManager Text User Interface, see the following section.
23. If the address of the new VM is automatically assigned, use Hyper-V Manager to identify the IP address. To assign a static IP to a VM, use the NetworkManager Text User Interface (nmtui) tool. For instructions, see [“Assigning a static IP address” on page 23](#).

What to do next

After you install the vSnap server, complete the following action:

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See “Managing vSnap servers” on page 40 .

Uninstalling a vSnap server

You can remove a vSnap server from your IBM Spectrum Protect Plus environment.

Before you begin

When permanently deleting the vSnap server, you must clean up the IBM Spectrum Protect Plus server. Items that must be cleaned up in this case, are as follows:

- Records of backups that are stored on the vSnap server.
- Replication relationships to other vSnap servers.
- Ensure that no jobs use SLA policies that define the vSnap server as a backup location.

To view the SLA policies that are associated with jobs, see the **Backup** page for the hypervisor or application that is scheduled for backup. For example, for VMware backup jobs, click **Manage Protection > Virtualized Systems > VMware**. You must unregister the vSnap server from the IBM Spectrum Protect Plus server. See [“Unregistering a vSnap server” on page 43](#) for more information.



Attention: Uninstalling a vSnap server can result in loss of data.

Procedure

1. Log on to the vSnap server console with the user ID `serveradmin`. The initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).

You can also use a user ID that has vSnap administrator privileges that you create by using the **vsnap user create** command. For more information about using console commands, see [“vSnap server administration reference” on page 60](#).

2. Run the following commands:

```
$ systemctl stop vsnap
$ yum remove vsnap
```

3. Optional: If you do not plan to reinstall the vSnap server after it is uninstalled, remove the data and configuration by running the following commands:

```
$ rm -rf /etc/vsnap
$ rm -rf /etc/nginx
$ rm -rf /etc/uwsgi.d
$ rm -f /etc/uwsgi.ini
```

4. Reboot the system to ensure kernel modules are unloaded and detach the data disks containing vSnap pool data.

Note: To uninstall IBM Spectrum Protect Plus in a Hyper-V environment, delete the IBM Spectrum Protect Plus appliance from Hyper-V and then delete the installation directory.

Results

After a vSnap server is uninstalled, the configuration is retained in the `/etc/vsnap` directory. The configuration is reused if the vSnap server is reinstalled. The configuration is removed if you ran the optional commands to remove the configuration data.

Managing vSnap servers

Each vSnap server is a stand-alone appliance, which is deployed virtually or installed physically on a system that meets the minimum requirements. Each vSnap server in the environment must be registered in IBM Spectrum Protect Plus so that the server is recognized.

Registering a vSnap server as a backup storage provider

Any vSnap server that is deployed virtually or installed physically must be registered in IBM Spectrum Protect Plus so that it can be recognized as a backup storage provider.

Before you begin

After you add and register a vSnap server as a backup storage provider, you can choose to configure and administer certain aspects of the vSnap, such as network configuration or storage pool management. For more information, see [“Configuring backup storage options”](#) on page 44.

If the vSnap server will also be registered as a VADP proxy, the account added in the **Storage Properties** field for the vSnap must have **sudo** privileges for the VADP proxy registration to succeed. For more information, see [“Permission types”](#) on page 463.

Procedure

To register a vSnap server as a backup storage device, complete the following steps:

1. Log on to the vSnap server console with the user ID `serveradmin`. The initial password is `sppDP758-SysXyz`.
You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus”](#) on page 96.
2. Run the **`vsnap user create`** command to create a user name and password for the vSnap server.
3. Start the IBM Spectrum Protect Plus user interface by entering the host name or IP address of the virtual machine where IBM Spectrum Protect Plus is deployed in a supported browser.
4. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
5. Click **Add vSnap server**.
The **Add vSnap server** wizard opens.
6. Complete the fields on the **vSnap server details** page, and then click **Next**.

Hostname/IP

Enter the resolvable IP address or hostname of the backup storage.

Note: The hostname or IP of the vSnap server as entered in the IBM Spectrum Protect Plus UI must exactly match one of the Subject Alternative Names (SANs) embedded in the vSnap certificate. Refer to [Certificate management](#) for detailed information on how to obtain and customize the vSnap certificate.

If you plan to protect Kubernetes or Red Hat OpenShift container workloads, you must enter the IP address or fully qualified domain name (FQDN) of the backup storage.

Site

Select a site for the backup storage. Available options are **Primary**, **Secondary**, or **Add a new site**. If more than one primary, secondary, or user-defined site is available to IBM Spectrum Protect Plus, the site with the largest amount of available storage is used first.

Use existing user

Enable this option to select the user name for the vSnap server that you created in step [“2”](#) on page 40.

If you do not select this option, complete the following fields to add a user:

Username

Enter a user name for the vSnap server.

Password

Enter a password for the user.

Certificate

In the **Certificate** field, select one of the following options to import the certificate:

Upload

- a. Download the `/etc/vsnap/ssl/spp-vsnap.crt` file from the vSnap server to the local machine where you are running the browser.
- b. Click **Choose file** and search for the downloaded certificate in your system.
- c. Click **Upload**.

Copy and paste

Enter a name for the certificate, such as `spp-vsnap.crt`. Then, paste the contents of the certificate in the **Copy and paste certificate here** field and click **Create**.

Use existing certificate

Click **Choose file** to select an existing certificate from the **Select a certificate** list. This option is the default.

Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that they key fingerprint matches the key fingerprint of the host that you are adding.

Requirement: If the `serveradmin` account is to be used, ensure that the default password is changed through the vSnap server console prior to registering the vSnap server as a backup storage provider in IBM Spectrum Protect Plus.

7. Review your selections, and then click **Submit.**

IBM Spectrum Protect Plus confirms a network connection and adds the backup storage device to the database.

What to do next

After you add a backup storage provider, take the following actions:

Action	How to
Initialize the vSnap server.	See “Initializing the vSnap server” on page 53 .
Expand the vSnap storage pool.	See “Configuring backup storage partners” on page 47 .

Action	How to
If necessary, configure and administer certain aspects of vSnap, such as network configuration or storage pool management.	See “Configuring backup storage options” on page 44

Related tasks

“Start IBM Spectrum Protect Plus” on page 96

Start IBM Spectrum Protect Plus to begin using the application and its features.

Editing settings for a vSnap server

You can edit the configuration settings for a vSnap server to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit the settings for a vSnap server, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server, and then click **Edit**.

Certificate

In the **Certificate** field, select one of the following options to import the certificate:

Upload

- a. Download the `spp-vsnap.crt` file from the vSnap server, located under `/etc/vsnap/ssl` to your local machine where you are running the browser on.
- b. Click **Choose file** and search for the downloaded certificate in your system.
- c. Click **Upload**.

Copy and paste

Enter a name for the certificate, such as `spp-vsnap.crt`. Then, paste the contents of the certificate that you exported into the **Copy and paste certificate here** field.

Click **Create**. After the certificate is created, click **Save**

Use existing certificate

It is a default option.

3. Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that they key fingerprint matches the key fingerprint of the host that you are adding.

4. Revise the vSnap server settings, and then click **Save**.

Unregistering a vSnap server

If required, you can unregister a vSnap server that is no longer used in your IBM Spectrum Protect Plus environment.

Before you begin

When a vSnap server is unregistering, all recovery points that are associated with the vSnap server are purged from IBM Spectrum Protect Plus during the next maintenance job.



Attention: Unregistering of a vSnap server can result in loss of data.

Before you unregister a vSnap server, review the scenarios to determine whether unregistering is appropriate or whether other action must be taken.

Scenario 1: The vSnap server is temporarily down due to storage or network issues.

- Do not unregister the vSnap server. If you unregister the vSnap server, recovery points that are associated with the server will be purged and backups will be rebased.
- Complete the necessary storage or network maintenance to bring the vSnap server back online.

Scenario 2: The vSnap server is assigned a new host name or IP address.

- Do not unregister the vSnap server. If you unregister the vSnap server, recovery points that are associated with the server will be purged and backups will be rebased.
- Edit the settings for the vSnap server to specify the new host name or IP address. To edit the settings for a vSnap server, follow the instructions [“Editing settings for a vSnap server”](#) on page 42.

Scenario 3: The vSnap server is not in use, and there are no plans to reuse it.

- Unregister the vSnap server and run a maintenance job to ensure that recovery points that are associated with the vSnap server are purged from IBM Spectrum Protect Plus.
 - Incremental backups of the data that was present on the vSnap server will no longer be possible.
 - Recovering data that was present on the vSnap server will no longer be possible.
- Subsequent runs of backup jobs will automatically create new volumes on another vSnap server in the same site and will perform new base backups.

Scenario 4: The vSnap pool is lost and you want to build a new pool on the same vSnap server.

1. Unregister the vSnap server and run a maintenance job to ensure that recovery points that are associated with the old vSnap pool are purged from IBM Spectrum Protect Plus.
 - Incremental backups of the data that was present in the old pool will no longer be possible.
 - Recovering data that was present in the old pool will no longer be possible.
2. On the vSnap server, create a pool.
3. Add the vSnap server back into IBM Spectrum Protect Plus. To add a vSnap server to IBM Spectrum Protect Plus, see [“Registering a vSnap server as a backup storage provider”](#) on page 40.
 - Subsequent runs of backup jobs will automatically create volumes on this or another vSnap server in the same site and will perform new base backups.

Scenario 5: The vSnap pool or server is lost and you intend to repair it. This can be achieved by replicating data from a vSnap replication server.

- Do not unregister the vSnap server from IBM Spectrum Protect Plus. The deletion process will cause backups to be rebased.
- Replace the vSnap server. For information about replacing a failed, primary vSnap server, see this section [“Troubleshooting vSnap servers”](#) on page 67.

Procedure

To unregister a vSnap server, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server, and then click **Delete device**.
3. Confirm removal of the vSnap server by entering the code in the text box. Click **UNREGISTER** to delete the server from IBM Spectrum Protect Plus.

Configuring backup storage options

You can configure additional storage-related options for your primary and secondary backup storage hosts.

Procedure

To configure backup storage options for registered disks, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Storage efficiency** tab.
4. Specify one or more storage options:

The screenshot shows the 'Storage efficiency' configuration page. It has a navigation bar with tabs: 'Storage efficiency' (selected), 'Storage Partners', 'Active Directory', 'Storage disk', 'Network interface controllers', and 'Streaming, retrieval, and allocation rules'. Below the navigation bar is a section titled 'Set Advanced Options' with three checked checkboxes: 'Enable compression', 'Enable deduplication', and 'Enable encryption'. Each checkbox has a 'Learn more' link. Below these options is a blue information box with an 'i' icon and the text: 'The Encryption setting can only be applied during vSnap initialization. This option is for informational purposes only.' At the bottom of the panel are two buttons: 'Cancel' and 'Save'.

Enable Compression: Select this option to compress each incoming block of data by using a compression algorithm before the data is written to the storage pool. Compression consumes a moderate amount of additional CPU resources.

Enable Deduplication: Select this option so that each incoming block of data is hashed and compared against existing blocks in the storage pool. If compression is enabled, the data is compared after it is compressed. Duplicate blocks are skipped instead of being written to the pool. Deduplication is deselected by default because it consumes a large amount of memory resources (proportional to the amount of data in the pool) to maintain the deduplication table of block hashes.

Encryption Enabled: This option displays the encryption status of the primary or secondary backup storage host. Encryption can be enabled only during vSnap initialization. This option cannot be changed in this pane.

5. Click **Save**.

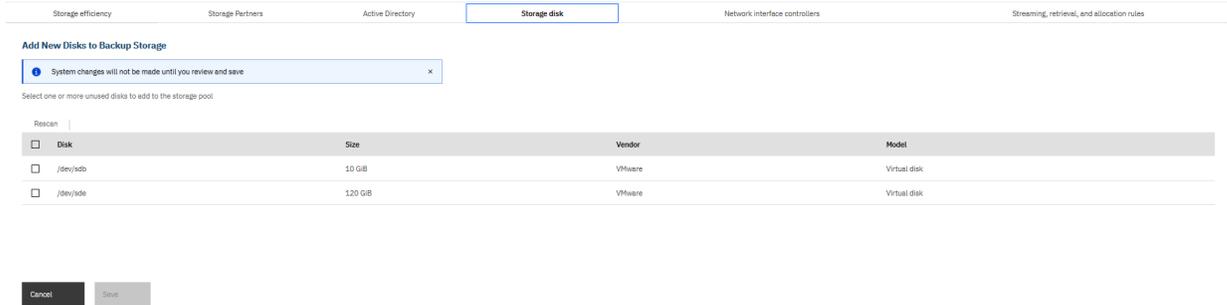
Adding new disks to backup storage

If you require more space for backup operations in a selected storage pool, you can add unused disk storage.

Procedure

To add new unused disks to a disk storage pool, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Storage disk** tab.
4. Select a disk to add to your storage environment from the list of available disks.



5. Click **Save**.

Rescanning a vSnap server after the storage is expanded

If you recently expanded the vSnap server storage pool by adding physical or virtual disks, you can rescan the vSnap server to pick up the additions. The operation rescans the entire vSnap server to pick up any recent storage pool additions.

Procedure

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to refresh, and then click **Manage**.
3. Open the **Storage disk** tab.
4. Click **Rescan** to scan the vSnap server for any storage pool expansion or changes.

This operation can several minutes to finish. The disk remains fully operational during the scanning process.

5. Optional: Select **Refresh** to refresh the details of the disk. For example, if the **Status/Capacity** figure has changed due to usage, the update is refreshed in the table.

Refreshing the disk storage for a vSnap server

You can refresh the disk storage view for your vSnap servers to show up-to-date status and capacity usage.

Procedure

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to refresh, and then click **Refresh**.

The information that is shown for the vSnap server is updated to reflect any changes. For example, if the **Status/Capacity** percentage changed due to usage or because you recently expanded the storage pool, the information is refreshed.

Configuring network interface controllers

You can configure your primary and secondary backup storage to use multiple network interface controllers (NICs) for different specific functions. The NICs in your IBM Spectrum Protect Plus environment can be configured to transfer data for backup, restore, and replication operations. You can configure a NIC for backup, restore, and replication data transfers, or for either backup and restore or replication data transfers. When you configure separate NICs, you can dedicate one network to replication operations and another network to backup and restore operations.

Before you begin

Versions of the vSnap server prior to version 10.1.6 do not support this feature. To update a vSnap server, follow the instructions in [“Updating vSnap servers” on page 88](#).

About this task

The network that is dedicated to sending management commands from IBM Spectrum Protect Plus to the vSnap server is indicated by the information icon . When you hover over the icon, **Management Network** is shown.

Connections can be established between the vSnap server and a range of clients, including application servers, hypervisor hosts, VADP proxies, and any other component in your environment that transfers data to and from backup storage.

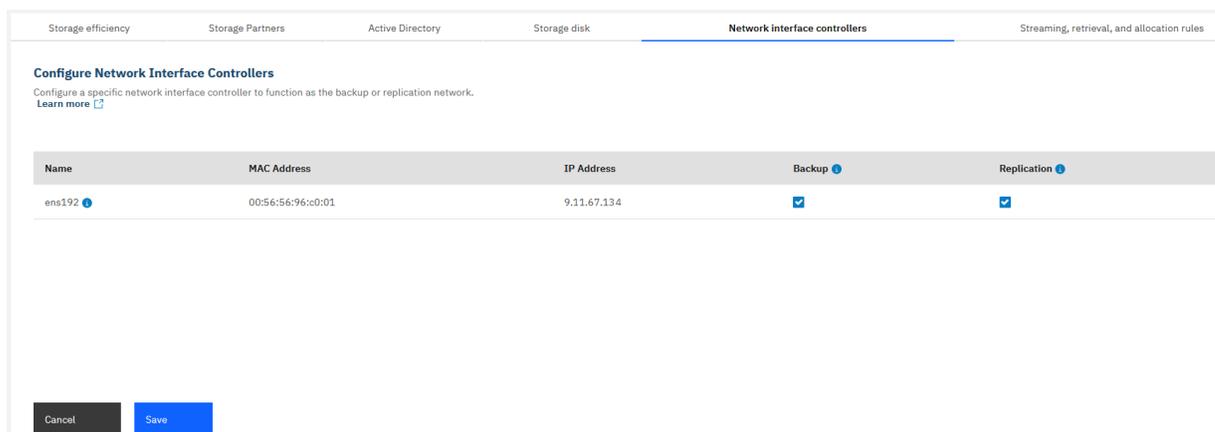
In the case that a second network interface card (NIC) is used, export `RMQ_SERVER_HOST` in the `/home/virgo/.bashrc` file on the IBM Spectrum Protect Plus appliance. The provided IP will be used for VADP to connect back to IBM Spectrum Protect Plus. The `<ip_address>` variable is the IP assigned to the second network interface card:

```
$ export RMQ_SERVER_HOST=<ip_address>
```

Procedure

To configure a NIC for backup and replication operations, complete the following steps:

1. In the navigation panel, click **System Configuration** > **Storage** > **vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Network interface controllers** tab.
4. Select the configuration that you want for your listed NICs:
 - To configure an NIC for transfers of data for backup and restore operations only, select **Backup**. During backup and restore operations, connections are established to the vSnap server by using the IP address of this NIC. If the **Backup** option is specified by multiple NICs, the first one that connects successfully is used.
 - To configure an NIC for transfers of data for replication purposes only, select **Replication**. During incoming replication operations to a vSnap server, connections are established using the IP address of this NIC on the target vSnap server. If the **Replication** option is specified for multiple NICs on the target vSnap server, the first target IP address that connects successfully from the source vSnap server is used.
 - To configure a NIC for both replication, and backup and restore data transfers, select both **Backup** and **Replication**.



5. Click **Save**.

Configuring backup storage partners

You can configure your backup storage primary and secondary sites to establish replication partnerships with other sites to extend your environment. After you configure replication partners, you can copy data from one site to another for an added layer of data protection.

Before you begin

All vSnap servers must be at the same version level for replication to function. Replication between different versions is not supported.

Procedure

To add partners to a server in your storage environment, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Storage Partners** tab.
4. Click **Add** to add a partner to your primary or secondary backup storage host.

Storage efficiency	Storage Partners	Active Directory	Storage disk	Network interface controllers	Streaming, retrieval, and allocation rules																
Configure Storage Partners																					
Current partners																					
<table border="1"><thead><tr><th>Hostname/IP</th><th>Created</th><th></th></tr></thead><tbody><tr><td>knob2.tucson.ibm.com</td><td>Jul 20, 2020 12:03:36 PM</td><td>Remove </td></tr></tbody></table>						Hostname/IP	Created		knob2.tucson.ibm.com	Jul 20, 2020 12:03:36 PM	Remove										
Hostname/IP	Created																				
knob2.tucson.ibm.com	Jul 20, 2020 12:03:36 PM	Remove																			
Available partners You can add the hosts below to current partners. Only compatible hosts are shown																					
<table border="1"><thead><tr><th>Hostname/IP and Pool</th><th>Capacity</th><th>Efficiency Settings</th><th></th></tr></thead><tbody><tr><td>knob9.tucson.ibm.com Site: Replication</td><td>Capacity: available of </td><td><input checked="" type="checkbox"/> Compression <input checked="" type="checkbox"/> Deduplication <input checked="" type="checkbox"/> Encryption</td><td>Add </td></tr><tr><td>knob7.tucson.ibm.com Site: Primary</td><td>Capacity: available of </td><td><input type="checkbox"/> Compression <input type="checkbox"/> Deduplication <input type="checkbox"/> Encryption</td><td>Add </td></tr><tr><td>usc-18.tucson.ibm.com Site: Primary</td><td>Capacity: available of </td><td><input checked="" type="checkbox"/> Compression <input type="checkbox"/> Deduplication</td><td>Add </td></tr></tbody></table>						Hostname/IP and Pool	Capacity	Efficiency Settings		knob9.tucson.ibm.com Site: Replication	Capacity: available of	<input checked="" type="checkbox"/> Compression <input checked="" type="checkbox"/> Deduplication <input checked="" type="checkbox"/> Encryption	Add	knob7.tucson.ibm.com Site: Primary	Capacity: available of	<input type="checkbox"/> Compression <input type="checkbox"/> Deduplication <input type="checkbox"/> Encryption	Add	usc-18.tucson.ibm.com Site: Primary	Capacity: available of	<input checked="" type="checkbox"/> Compression <input type="checkbox"/> Deduplication	Add
Hostname/IP and Pool	Capacity	Efficiency Settings																			
knob9.tucson.ibm.com Site: Replication	Capacity: available of	<input checked="" type="checkbox"/> Compression <input checked="" type="checkbox"/> Deduplication <input checked="" type="checkbox"/> Encryption	Add																		
knob7.tucson.ibm.com Site: Primary	Capacity: available of	<input type="checkbox"/> Compression <input type="checkbox"/> Deduplication <input type="checkbox"/> Encryption	Add																		
usc-18.tucson.ibm.com Site: Primary	Capacity: available of	<input checked="" type="checkbox"/> Compression <input type="checkbox"/> Deduplication	Add																		

Configuring an Active Directory

You can associate your primary and secondary backup storage with an active directory domain. When the primary or secondary host is added to a domain, any Microsoft SQL Server log backup jobs that are associated with that host will use domain authentication to mount the log backup volume. In this way, you can avoid the requirement to use a local staging area on the application server when for log backup operations.

Before you begin

You might have to configure the Domain Name System (DNS) server so that the domain controller is available to the network and can be associated with the primary or secondary host.

Procedure

To add an Active Directory for backup and restore operations, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Active Directory** tab.
4. Enter the domain name of the Active Directory, along with the user name and password for the Active Directory administrator as shown in the following figure.

5. Click **Join**.

Configuring advanced storage options

You can set advanced storage-related options for the primary or secondary backup storage in your environment.

Procedure

To configure advanced options for your backup storage, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Storage and Transport Options** tab.
4. Configure the advanced options.

Set Storage Options

Concurrent stream limit for copy to archive object storage	5	-	+
Concurrent stream limit for copy to standard object storage	5	-	+
Concurrent stream limit for replication	5	-	+
Enable Transport Encryption (has additional requirements, see documentation)	<input checked="" type="checkbox"/>		
Interval in seconds between volume/snapshot deletions during space reclamation	300	-	+
Retrieval tier for restore from AWS archive object storage (Bulk, Standard, or Expedited)	Bulk		

Figure 4. Manage backup storage advanced options.

- **Concurrent stream limit for copy to archive object storage:** This value defines the maximum number of concurrent streams that are used by this backup host when you are copying data to archive Object Storage.
- **Concurrent stream limit for copy to standard object storage:** This value defines the maximum number of concurrent streams that are used by this backup host when you are copying data to standard Object Storage.

- **Concurrent stream limit for replication:** This value defines the maximum number of concurrent streams that are used by this backup host when you are replicating data to other backup hosts.
- **Enable Transport encryption:** Select this option to enable encryption of backup data while it is transferred to or from the vSnap Server. For more information about transport encryption, see [“Transport encryption” on page 50](#).

Note: By default, this option is disabled to preserve the legacy behavior and also to ensure that backup and restores can continue seamlessly when updating vSnaps.

- **Interval in seconds between volume/snapshot deletions during space reclamation:** This value defines the interval in seconds between successive deletions of volumes or snapshots on the vSnap server when space is reclaimed following a run of Maintenance jobs. Lowering the interval allows space to be reclaimed more aggressively, particularly when a large amount of data has expired in bulk.

Important: Aggressive reclamation can put input and output load on the vSnap pool which can result in slower performance for other concurrent workloads.

- **Retrieval tier for restore from AWS archive object storage (Bulk, Standard, or Expedited):** This value specifies the retrieval tier that is used by this backup host during restore operations from Amazon Glacier archive Object Storage. This value must be specified as Bulk, Standard, or Expedited. The retrieval tier can be modified to achieve faster restore operation times at the cost of higher data charges. For information about the available retrieval tier options and associated pricing, see the Amazon Web Services documentation.
- **Concurrent Backup:** This option specifies the maximum number of parallel backup streams to the host when multiple jobs that run concurrently. For application backup operations, each database is treated as a single stream. For hypervisor backup operations, each virtual disk is treated as a single stream. The concurrent backup options can be used to prevent multiple or large SLA policies from sending too many data streams to a small backup host that cannot accommodate the load. To reduce processing time for backup operations, set this option to one of the following options:

Unlimited: an unlimited number of concurrent backup streams can run.

Pause: to pause the use of this backup host. Jobs attempting to utilize this backup host will pause while this setting is selected. This option should be used in situations where the backup host requires emergency maintenance and will temporarily prevent it from being used by any jobs.

Limit: to set a maximum limit on the number of backup streams that can run concurrently. Enter a numerical value specifying the maximum number of concurrent streams.

- **Disable New Application:** Enabling this option will make it so that the vSnap server will not be used for new storage application for virtual machine (VM) backups. Existing virtual machine backups will continue using the vSnap server. If a virtual machine backup requires new storage, it will not use the vSnap regardless of the remaining free space or VM allocation setting for the assigned site.

Tip: When you change an option value, the new value is applied when you click into the next option field. Alongside the updated option, the following message is displayed,  Updated .

5. Click **Close**.

Related reference

[“Transport encryption” on page 50](#)

IBM Spectrum Protect Plus 10.1.13 introduces **Transport Encryption** feature to protect the data transport between application host and vSnap during backup and restore. With the transport encryption, each data path of data between the application host and the vSnap can be encrypted and decrypted.

Transport encryption

IBM Spectrum Protect Plus 10.1.13 introduces **Transport Encryption** feature to protect the data transport between application host and vSnap during backup and restore. With the transport encryption, each data path of data between the application host and the vSnap can be encrypted and decrypted.

Considerations to use transport encryption

To enable transport encryption, ensure that the prerequisite software is at the required level and all security-related patches are applied. For system requirements, see [technote 6837823](#).

Important:

- If you are using IBM Spectrum Protect Plus for backup storage and want to protect the data transport with transport encryption option, you must update both IBM Spectrum Protect Plus and vSnap to 10.1.13 or later releases.
- After installing or updating to IBM Spectrum Protect Plus and vSnap to 10.1.13 or later, the transport encryption option is disabled by default. To enable the transport encryption option, see [“Configuring advanced storage options”](#) on page 48.
- After you enable transport encryption in IBM Spectrum Protect Plus 10.1.13 or later and plan to disable it, you must manually disable the transport encryption option.

Review the following information before you enable transport encryption:

- When you enable the transport encryption, each data stream of data between the application host and the vSnap will be encrypted and decrypted. Each stream is handled by one CPU core. Data transport encryption can increase CPU usage, which can affect the system performance. The potential impact on performance depends on CPU types, number of vSnaps, hosts involved in an service level agreement (SLA) and various other factors. The performance may reduce 10% to 50% depending on data types and setup.
- You can fully protect the following data types:
 - SQL database and log backups
 - Exchange database and log backups
 - Windows file system
 - Oracle database and log backups on the Linux[®] systems
 - Db2 database and Log backups on the Linux[®] systems
 - MongoDB

Note: MongoDB does not have log backup.

- You can partially protect the following data types:
 - SAP HANA: You can enable transport encryption feature for SAP HANA DB. Due to technical limitations, you cannot protect SAP HANA log backups with transport encryption. To protect your SAP HANA log backup data, you must enable SAP HANA backup encryption.

For more information, see [“Enabling log encryption for SAP HANA data”](#) on page 413.

- VMware: You can protect the data transport between the vSnap and a remote VADP with the IBM Spectrum Protect Plus transport encryption feature. Also, the path is always protected when you back up VMware data to Open Snap Store Manager (OSSM). When you backup VMware data, the VADP reads the data from the data store and sends it to vSnap. You cannot enable IBM Spectrum Protect Plus transport encryption to the data store connection.

For more information about how to enable transport encryption on VMware to protect VMware data, see [“Enabling transport encryption for VMware data”](#) on page 242.

- Due to technical limitations, you cannot protect the following data types:
 - Hyper-V
 - Oracle database and log backups on the AIX® systems
 - Db2 database and log backups on the AIX® systems
 - SAP HANA database and log backups on the AIX® systems
 - Microsoft 365

How do I delete and recreate a vSnap storage pool?

When a scenario arises that results in the requirement to delete a vSnap storage pool due to corruption or any other reason, you can follow the steps to delete and recreate the storage pool. This procedure is a destructive operation that discards all data in an existing vSnap storage pool. All backup data in the pool is lost, and is no longer recoverable so caution is needed before you proceed. After that is done, you can create a replacement empty pool.

Procedure

1. To prepare for the removal of a storage pool, you must first unregister the vSnap server by removing it.

For more information about unregistering the vSnap server, see [“Unregistering a vSnap server”](#) on page 43.

2. Run a maintenance job on the vSnap server by opening **Jobs and Operations > Schedule**. Select a job in the list, and then click **Start**.

When the maintenance job is completed, all the information about the vSnap server is removed from the IBM Spectrum Protect Plus catalog. All recovery points and metadata that are associated with the VM backups, and all replica copies that are stored in the unregistered vSnap, are removed. All data is removed and is no longer available for recovery.

For more information about maintenance jobs, see [“Job types”](#) on page 431.

3. On the vSnap server, run the following command to initialize the cleaned vSnap server.

```
$ vsnap system init --skip_pool
```

If the system was initialized previously, it is safe to run this command again. This step ensures that required kernel modules are installed and loaded.

4. Identify the existing storage pool identifier by running the following command:

```
$ vsnap pool show
```

If the storage pool is online, the identifier is displayed in the *ID* field. If the storage pool is offline, an error message displays that indicates the pool information cannot be displayed. The identifier of the pool is shown in this error message.

5. Run the delete command for the storage pool identifier to forcibly delete the storage pool.

```
$ vsnap pool delete --id <ID> --force
```

When the command is finished, the following message is displayed:

```
Storage pool was deleted successfully but the pool was not unmounted because the 'force'
option was set.
Reboot the system to ensure disks that were previously in use are released.
```

6. Restart the system to release any disks that are still in use. Enter the following command:

```
$ sudo reboot -n
```

It is important to restart the system after you run this command to ensure that any disks that are still in use by older pools are released.

- When the restart finishes, run the status command:

```
$ vsnap_status
```

This output of this command shows the status of all vSnap server services. Ensure that all services are active. If one or more services are activating, check the status later until they are all in the active state.

- Identify the disks that must be added to the pool.

If you are reusing the same set of disks that comprised the old pool, the following command can help you to identify them:

```
$ vsnap disk show
```

In the output of the show command, the **USED AS** column indicates whether a file system or partition table exists on the disk. Disks that were part of the old pool are identified as `vsnap_pool1`. If the old pool was encrypted, some or all disks can be identified as `crypto_LUKS`.

Sample output

UUID KNAME NAME	TYPE	VENDOR	MODEL	SIZE	USED AS
6000c299371bdcb647c80720602079bc sda /dev/sda	SCSI	VMware	Virtual disk	70.00GB	LVM2_member
6000c29b8ea25349e3a884d58f72e640 sdb /dev/sdb	SCSI	VMware	Virtual disk	100.00GB	vsnap_pool
6000c297cb8078cf9f56ab688a326a24 sdc /dev/sdc	SCSI	VMware	Virtual disk	128.00GB	LVM2_member
6000c2950248c5d831b6661ab0ec8843 sdd /dev/sdd	SCSI	VMware	Virtual disk	16.00GB	vsnap_pool
6000c29359661cbd915a7f24c8b44cf8 sde /dev/sde	SCSI	VMware	Virtual disk	16.00GB	vsnap_pool

- Important:** The command in this step deletes partition tables and file system metadata from the specified disks, and marks them as unused. Use this command with caution, and ensure that you specify only disks that are no longer in use.

Run the following command to specify a comma-separated list of disk names to mark as unused.

```
$ vsnap disk wipe <disk_list>
```

The following command is an example of the disk wipe command: `$ vsnap disk wipe /dev/sdb,/dev/sdd,/dev/sde`.

- Create the new pool with the following command:

```
$ vsnap pool create --name <pool_name> <options> --disk_list <disk_list>
```

Where *pool_name* is the name of the new pool; *options* specifies RAID type or encryption options. Leaving this option blank applies the default options. *disk_list* represents the comma-separated list of disks to be added to the pool. The disks that you specify must have a status of unused when you run the **vsnap disk show** command.

The following command is an example of the create command:

```
$ vsnap pool create --name primary --disk_list /dev/sdb,/dev/sdd
```

When you are specifying the list of disks, specify only the disks that you intend to use as the main data disks. Cache or log disks can be added later by running separate commands. For more information about recommendations and instructions for configuring cache and log disks, see the [Blueprints](#).

Tip:

To open help, run the `vsnap pool create --help` command.

11. To view the pool information, run the following command:

```
$ vsnap pool show
```

Ensure that the command displays the correct pool information and that the command completes without an error.

12. Register the vSnap server in IBM Spectrum Protect Plus under a chosen site to finalize the setup. For more information about how to register a vSnap server, see [“Registering a vSnap server as a backup storage provider”](#) on page 40.

Initializing the vSnap server

The initialization process prepares a new vSnap server for use by loading and configuring software components and initializing the internal configuration. This is a one-time process that must be run for new installations.

About this task

During the initialization process, vSnap creates a storage pool using any available unused disks attached to the system for a physical installation. If no unused disks are found, the initialization process completes without creating a pool. For a virtual deployment of vSnap, a default 100 GB unused virtual disk is defined and used to create the pool.

For information about how to expand, create, and administer storage pools, see [“Storage management”](#) on page 62.

You can use the IBM Spectrum Protect Plus user interface or the vSnap command line interface (CLI) to initialize vSnap servers.

For servers that are deployed and added to IBM Spectrum Protect Plus, the IBM Spectrum Protect Plus user interface provides a simple method to run the initialization operation.

For servers that are deployed in a physical environment, the vSnap command line interface (CLI) offers more options for initializing the server, including the ability to create a storage pool by using advanced redundancy options and a specific list of disks.

Completing a simple initialization

To prepare a vSnap server for use, you must initialize the vSnap server. Use the IBM Spectrum Protect Plus to initialize a vSnap server that is deployed in a virtual environment.

Procedure

To initialize a vSnap server by using the IBM Spectrum Protect Plus user interface, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to initialize, and then click one of the following options:

Initialize

Initialize the vSnap server without encryption enabled.

Initialize with Encryption

Enable encryption of backup data on the vSnap server.

The initialization process runs in the background and requires no further user interaction. The process might take 5 - 10 minutes to complete.

Completing an advanced initialization

Use the vSnap server console to initialize a vSnap server that is deployed in your environment. Initializing by using the vSnap server console offers more options for initializing the server, including the ability to create a storage pool by using advanced redundancy options and a specific list of disks.

Procedure

To initialize a vSnap server by using the vSnap server console, complete the following steps:

1. Log in to the vSnap server console with the user ID `serveradmin` by using SSH. When deployed virtually, the initial password is `sppDP758-SysXyz`. You will be prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#). If deployed physically, use the password that you created for the `serveradmin` account during installation.

You can also use a user ID that has vSnap privileges that was previously created using the **`vsnap user create`** command. For more information about using console commands, see [“vSnap server administration reference” on page 60](#).

2. Issue the **`$ vsnap system init`** command with the **`--skip_pool`** option to initialize the vSnap server without creating a storage pool. The process might take 5 - 10 minutes to complete. Issue the following command:

```
$ vsnap system init --skip_pool
```

What to do next

After you complete the initialization, complete the following action:

Action	How to
Create a storage pool	See “Storage management” on page 62 .

Migrating onboard vSnap data to a stand-alone vSnap server

Beginning with IBM Spectrum Protect Plus Version 10.1.7, the onboard vSnap server is no longer included. If you upgrade your system to IBM Spectrum Protect Plus 10.1.7 or later, but data remains in an onboard vSnap server from a previous release prior to version 10.1.7, you must migrate the data to a new, stand-alone vSnap server.

Before you begin

Beginning with IBM Spectrum Protect Plus 10.1.7, new deployments will no longer contain an onboard vSnap server. Systems upgraded from a previous version of IBM Spectrum Protect Plus still contain an onboard vSnap server which can be part of the Demo site. The onboard vSnap server will no longer be upgraded as part of general updates to IBM Spectrum Protect Plus.

The **`LocalvSnapAdmin`** identity was used as the identity to connect to the onboard vSnap server. In some cases, this identity may have been used to access other vSnap servers. If the identity was used to connect to other vSnap servers, a new identity for those servers must be created. Use the **`serveradmin`** account to connect to vSnap servers.

Do not unregister the onboard vSnap server from IBM Spectrum Protect Plus until prompted.

Ensure that sufficient space is available on the datastore for a stand-alone vSnap server deployment.

Do not explicitly initialize the new vSnap server that will be deployed as part of this procedure. Instead, the configuration of the onboard vSnap server will be copied to the new vSnap server.



Attention: Follow the procedure carefully. If not followed, this procedure can result in a loss of data.

About this task

In previous releases, an onboard vSnap server was included for proof-of-concept (POC) and demo purposes. The vSnap server was named localhost and was part of the Primary site by default. Beginning with IBM Spectrum Protect Plus 10.1.5, the onboard vSnap server was part of a Demo site that provided limited functionality. Users were able to manually remove the onboard vSnap from the Demo site and then register it with another site at which point the vSnap server was no longer limited in functionality.

Determine whether an onboard vSnap server was used in the previous release. Users who did not unregister the onboard vSnap from the Demo site will follow a different procedure from users who unregistered the onboard vSnap server from the Demo site and assigned the server to another site. Consider the two scenarios below:

Scenario 1: If the onboard vSnap was unused or previously used only in the Demo site, stop using the onboard vSnap. Unregister the vSnap from IBM Spectrum Protect Plus, for more information, see [“Unregistering a vSnap server” on page 43](#). After completing those steps, uninstall the vSnap software from the IBM Spectrum Protect Plus server. Skip the steps and begin with Step 9 in this procedure.

Scenario 2: If the onboard vSnap was unregistered from the Demo site and used in production under another site, do not unregister the onboard vSnap server from IBM Spectrum Protect Plus. The procedure in this topic will reference other topics. It may be helpful to have these topics open when:

- Manually upgrade the onboard vSnap server using the appropriate `.run` file. Follow to the general procedure for upgrading an external vSnap server as described in the [“Updating a vSnap server” on page 89](#) topic.
- Deploy a new stand-alone vSnap server using the most recent OVA. For more information, see [“Installing a virtual vSnap server in a VMware environment” on page 36](#).
- Upon completing the migration of data, uninstall the vSnap software from the IBM Spectrum Protect Plus server. These steps are detailed in this procedure beginning with Step 9.

Procedure

1. Update the onboard vSnap server and collect the vSnap pool information.
 - a) Using secure shell (SSH), log in to the onboard vSnap as the **serveradmin** user.
 - b) Upgrade the vSnap server to the most recent release. For more information, see [“Updating a vSnap server” on page 89](#).
 - c) Determine the version level of the vSnap server. At the command prompt, issue the following command:

```
$ vsnap system info
```
 - d) Determine all the disks labeled `vsnap_poo1`. The storage pool is comprised of these disks which will be detached from the onboard vSnap server and attached to the new vSnap server later in this procedure. At the command prompt, issue the following command to identify the disks:

```
$ vsnap disk show
```
2. Deploy a new, stand-alone vSnap server using the most recent `.ova`, apply custom settings, and verify the version level.
 - a) Log in to the vSphere Client.
 - b) Deploy a new stand-alone vSnap server using the most recent version of the vSnap `.ova`. For more information, see [“Installing a virtual vSnap server in a VMware environment” on page 36](#).
 - c) The new, stand-alone vSnap server will contain an unused 100GB disk that is used as the initial disk for creating a new storage pool. Detach this disk from the stand-alone vSnap server and delete it.
 - d) Configure the network properties as appropriate to your environment on the newly created vSnap server. Document the IP address or hostname for later use in this procedure.
 - e) Using secure shell (SSH), log in to the newly created vSnap as the **serveradmin** user.

- f) Determine the version level of the newly created vSnap server. At the command prompt enter the **vsnap system info** command:

```
$ vsnap system info
```

This version should match the version level of the onboard vSnap server that was upgraded and verified in the first step. If not, upgrade one or both of the vSnap servers to the latest release to ensure that they are at the same version level.

3. Pause all jobs in IBM Spectrum Protect Plus, document replication partnerships, and delete the partnerships from the onboard vSnap.
 - a) Log on to the IBM Spectrum Protect Plus server.
 - b) Jobs must not be actively running or scheduled to run during the migration procedure. Pause the schedule for all jobs to ensure that they do not attempt to run while the migration is occurring. Click **Jobs and Operations > Schedule** and then click **Pause All Jobs**. Verify that no jobs are running by clicking **Jobs and Operations > Running Jobs**.
 - c) Modify the settings for the onboard vSnap server. Navigate to **System Configuration > Storage > vSnap servers**, select the onboard vSnap server, and then click **Manage**.
 - d) Open the **Storage Partners** tab. Note the IP address or hostname of each replication partner for later use in this procedure.
 - e) Remove each replication partner. Removing the partnerships will not affect the replication data. The partnerships will be re-created in a subsequent step after the migration is complete.
4. Backup the onboard vSnap server configuration, transfer the configuration file to the new stand-alone vSnap server, and stop and disable the vSnap services on the onboard vSnap server.
 - a) Using secure shell (SSH), log in to the onboard vSnap server as the **serveradmin** user.
 - b) Create a backup of the vSnap configuration using the **vsnap system config backup** command. In this example, the config backup is saved in the root of the **serveradmin** user's home directory:

```
$ vsnap system config backup --outfile /home/serveradmin/vsnap_config_backup.tar.gz
```

- c) Copy the `vsnap_config_backup.tar.gz` from the onboard vSnap server to the newly created stand-alone vSnap server into the `/home/serveradmin` directory. SCP can be used to copy the file. In this example, `ip_address_new_vsnap` is a variable used to denote the IP address of the newly created stand-alone vSnap server. If prompted, accept the fingerprint and enter **yes** to continue connecting.

```
$ scp vsnap_config_backup.tar.gz serveradmin@ip_address_new_vsnap:/home/serveradmin
```

- d) Enter the password for the **serveradmin** account on the stand-alone vSnap server. The file will begin transferring.
 - e) Disable the vSnap services for the onboard vSnap server using the **systemctl stop** and **systemctl disable** commands:

```
$ sudo systemctl stop vsnap
```

```
$ sudo systemctl disable vsnap
```

5. Restore the onboard vSnap server configuration to the new stand-alone vSnap server.
 - a) Using secure shell (SSH), log in to the newly created vSnap as the **serveradmin** user.
 - b) Restore the config backup from the onboard vSnap server to the stand-alone vSnap server using the **vsnap system config restore** command:

```
$ vsnap system config restore --file /home/serveradmin/vsnap_config_backup.tar.gz
```

6. Power off the onboard vSnap server and the stand-alone vSnap server, detach the disks from the onboard vSnap and attach the disks to the stand-alone vSnap server. Power on both vSnap servers.
 - a) Log in to the vSphere Client.

- b) Power off the onboard vSnap and the stand-alone vSnap virtual machines and edit the settings of the virtual machine that has the onboard vSnap.
 - c) Detach the disks associated with the vSnap pool that is to be migrated as identified in Step 1d.
 - d) Edit the settings of the stand-alone vSnap virtual machine and attach the disks that were detached from the onboard vSnap server in Step 6c.
 - e) Power on the onboard vSnap and the stand-alone vSnap virtual machines.
7. Verify the status of both the onboard vSnap server and the newly deployed stand-alone vSnap server.
- a) Using secure shell (SSH), log in to the onboard vSnap server as the **serveradmin** user.
 - b) Run the **vsnap_status** command to determine the status of the vSnap services on the onboard vSnap server. It is expected that the services will no longer be running since the **systemctl stop** and **systemctl disable** commands were previously executed in Step 4.

```
$ vsnap_status
```

- c) Using secure shell (SSH), log in to the newly created vSnap as the **serveradmin** user.
- d) Run the **vsnap_status** command to determine the status of the vSnap services on the stand-alone vSnap server. The expected outcome is that the services will start and mount the storage pool.

```
$ vsnap_status
```

Note: It may take up to 15 minutes for all services to start. Periodically run the **vsnap_status** command to check the status.

- e) After all vSnap services are active, execute the **vsnap pool show** command to verify that the storage pool is online:

```
$ vsnap pool show
```

8. Update the vSnap server registration, the associated credentials, re-add the replication partners, and release the job schedules.
- a) Log on to the IBM Spectrum Protect Plus server.
 - b) Click **System Configuration > Storage > vSnap servers**, select the onboard vSnap server, and then click **Edit**.
 - c) Enter the IP address or the hostname in the **Hostname/IP** field of the newly created stand-alone vSnap server.
 - d) The existing user may display as **LocalvSnapAdmin** or as another identity. Deselect **Use existing user**. Enter **serveradmin** in the **User ID** field and the associated password for the stand-alone vSnap server in the **Password** field.
 - e) Click **Save**.
 - f) On the **vSnap servers** section, select the vSnap server that you edited and click **Refresh**.
 - g) After the refresh operation, verify that the information for the vSnap server is accurate.
 - h) Click **Manage** and then open the **Storage Partners** tab.
 - i) Re-enter the replication partners that were removed in Step 3. For instructions for entering partners, see [“Configuring backup storage partners”](#) on page 47.
 - j) Release schedules for all jobs that were paused in Step 3. Navigate to **Jobs and Operations > Schedule**, and then click **Release All Schedules**.
9. Remove the vSnap software from the IBM Spectrum Protect Plus server.
- a) Using secure shell (SSH), log in to the IBM Spectrum Protect Plus server as the **serveradmin** user.
 - b) Execute the **yum remove** commands to remove the vSnap server software from the IBM Spectrum Protect Plus server:

```
$ sudo yum remove vsnap
```

```
$ sudo yum remove vsnap-dist
```

Results

The migration from the onboard vSnap to a newly created stand-alone vSnap server is complete. All jobs that used the onboard vSnap will now use the new vSnap server. All data previously backed up to the onboard vSnap can be restored from the new vSnap server. Previously scheduled backup, replication, and cloud copy jobs will continue, as data is incrementally transferred to the new vSnap server.

Expanding a vSnap storage pool

If IBM Spectrum Protect Plus reports that a vSnap server is reaching its storage capacity, the vSnap storage pool must be expanded. To expand a vSnap storage pool, you must first add virtual or physical disks on the vSnap server. To add disks, choose to either add virtual disks to the vSnap virtual machine or add physical disks to the vSnap physical server.

Before you begin

Virtual or physical disks must be added to the vSnap server before you follow this procedure. Expanding existing volumes is not supported. See the vSphere documentation for information about creating new virtual disks.

Note:

Once a disk has been added to the storage pool, it cannot be removed. Detaching a disk that is in use by the pool can make the pool unusable.

Procedure

To expand a vSnap storage pool, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server that you want to configure, and then click **Manage**.
3. Open the **Storage disk** tab, and then click **Rescan**.

The rescan discovers newly attached disks on the vSnap server. When the rescan completes, any disks that are unpartitioned and unformatted, and therefore unused, are displayed in the list.

4. Select one or more disks from the list, and then click **Save**.

The selected disks are added to the vSnap storage pool, which expands the capacity of the vSnap pool by the size of the disks that are added.

What to do next

After you expand the storage pool, rescan the disk or vSnap server to pick up the new disks. For instructions on how to run a rescan operation, see [“Rescanning a vSnap server after the storage is expanded”](#) on page 45.

Changing the throughput rate

Change the throughput for site replication and copy operations so that you can manage your network activity on a defined schedule.

Procedure

1. In the navigation panel, click **System Configuration > Storage > Sites**.
2. Select the site that you want to configure, and then click **Edit**.
3. Click **Enabled** for **Throughput throttle** option.
4. Adjust the throughput:
 - Change the numerical rate of throughput by clicking the up and down arrows.

- Select a unit for the throughput. The default throughput is 100 MB per second.

Site details

Site name

Secondary

Throughput throttle

To manage the network activity on a defined schedule, select Enabled and change the throughput for site replication and copy operations.

Disabled

Enabled

Throttle rate

525 MB per second

Throttle schedule

Select times that the throttle is active.

Sunday from 7:00 to 7:59; Monday through Wednesday from 8:00 to 8:59; Thursday from 1:00 to 1:59, from 8:00 to 8:59; Friday from 8:00 to 8:59; Saturday from 4:00 to 4:59, from 8:00 to 8:59

VMware VM allocation

Cancel Save

Figure 5. Enabling different throttles for different times to improve throughput

5. Select times for the changed throughput in the weekly schedule table, or specify a day and time for the changed rate.

Tip: To clear a time slot, click the time slot. The scheduled selections are listed underneath the schedule table.

6. Click **Save** to commit the changes and close the panel.

Replacing a failed vSnap server

In an IBM Spectrum Protect Plus environment, the target vSnap server is the destination for backing up data. If the vSnap server becomes corrupted or fails to respond, you can replace the vSnap server with a new server and recover the stored data.

Before you begin

Important: Do not unregister the failed vSnap server from IBM Spectrum Protect Plus. The failed server must remain registered for the replacement procedure to work correctly.

One or more active, initialized vSnap replica servers must exist in the environment to successfully complete this process.

About this task

The procedure for replacing a failed vSnap server is documented in [technote 1103847](#).

Installing iSCSI initiator utilities

You must install Internet Small Computer System Interface (iSCSI) utilities if iSCSI mounted storage devices are directly connected to a vSnap server. After the iSCSI initiator utilities are installed, iSCSI mounted storage devices can be connected to the server on which the package is installed.

About this task

iSCSI initiator utilities can be installed on a vSnap server. The iSCSI initiator utilities are delivered with IBM Spectrum Protect Plus, but are not installed automatically. To install the utilities, complete the following steps:

Procedure

1. Log on to the vSnap server that is to be directly connected to the iSCSI mounted storage.
Use SSH or access the server directly and authenticate with the appropriate administrative credentials.
2. Install the iSCSI initiator utilities by running the following command:

```
sudo /usr/bin/yum --disablerepo=* --enablerepo=base,updates install iscsi-initiator-utils
```

vSnap server administration reference

After the vSnap server is installed, registered, and initialized, IBM Spectrum Protect Plus automatically manages its use as a backup target. Volumes and snapshots are created and managed automatically based on the SLA policies that are defined in IBM Spectrum Protect Plus.

You might have to configure and administer certain aspects of vSnap, such as network configuration or storage pool management.

Managing vSnap by using the command line interface

The vSnap server can be managed through the command-line interface and is the primary means of administering a vSnap server. Run the **vsnap** command from the vSnap server's interface after connecting through SSH using the user ID `serveradmin` or any other operating system user who has been assigned vSnap admin privileges. The initial `serveradmin` password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).

The command line interface consists of several commands and sub-commands that manage various aspects of the system. You can also pass the **--help** flag to any command or subcommand to view usage help, for example, **vsnap --help** or **vsnap pool create --help**.

Managing vSnap by using the IBM Spectrum Protect Plus user interface

Some common operations can be completed from the IBM Spectrum Protect Plus user interface. Log in to the user interface and click **System Configuration > Storage > vSnap servers** in the navigation panel. Select a vSnap server, and then click **Manage** to configure the server settings.

Related tasks

[“Managing vSnap servers” on page 40](#)

Each vSnap server is a stand-alone appliance, which is deployed virtually or installed physically on a system that meets the minimum requirements. Each vSnap server in the environment must be registered in IBM Spectrum Protect Plus so that the server is recognized.

[“Configuring advanced storage options” on page 48](#)

You can set advanced storage-related options for the primary or secondary backup storage in your environment.

User management

You can manage vSnap server users by issuing the **vsnap user** command. This command and available options are used to create users, grant and revoke user privileges, query users, and update a user's password.

Users that are created on a vSnap server are operating system users that are added to the vSnap operating system group. Users in the vSnap operating system group are not assigned **sudo** privileges. As a result, these users require a password to run a command.

You can create a vSnap user by issuing the **create** command. In this way, you create an operating system user that is assigned to the **vsnap** group that can run vSnap commands and make API calls. Issue the **create** command:

```
$ vsnap user create
```

If running interactively, you are prompted to enter the username, password, and the password a second time for confirmation. If running non-interactively, the following options are available to the **create** command:

--username <username>

Enter the username of the user.

--password <password>

Enter the password of the user.

You can grant privileges to an existing operating system account to ensure that the user can run vSnap commands and make API calls. To grant privileges, issue the **grant** command:

```
$ vsnap user grant
```

If running interactively, you are prompted to enter the username, password, and the password a second time for confirmation. If running non-interactively, the following options are available to the **grant** command:

--username <username>

Enter the username of the user.

--password <password>

Enter the password of the user. This must be the operating system account password if the account already exists on the system.

You can revoke privileges from a user who is assigned to the **vsnap** group. The user will remain as an operating system user but will no longer be able to run vSnap commands or make API calls. To revoke privileges, issue the **revoke** command:

```
$ vsnap user revoke
```

If running interactively, you are prompted to enter the username. If running non-interactively, the following options are available to the **revoke** command:

--username <username>

Enter the username of the user.

To display a list of vSnap users who are part of the **vsnap** group on the vSnap server, issue the **show** command:

```
$ vsnap user show
```

A vSnap user can have the account password changed which will update that user's password on the system. Issue the **update** command:

```
$ vsnap user update
```

If running interactively, you are prompted to enter the username, old password, new password, and the new password a second time for confirmation. If running non-interactively, the following options are available to the **update** command:

--username <username>

Enter the username of the user.

--password <old_password>

Enter the old password of the user.

--new_password <new_password>

Enter the new password of the user.

If you have already changed the vSnap system password using an external command instead of the `vsnap user update` command, then the SMB/CIFS password can be output of sync with the system password. Issue the following command to synchronize the SMB/CIFS passwords.

```
vsnap user resyncsmbpass
```

If running interactively, you are prompted to enter the `username` and the current system password. If running non-interactively, the following options are available to the **resyncsmbpass** command.

--username <username>

Enter the username of the account to be synced.

--password <password>

Enter the current system password of the account to be synced.

Storage management

You can create and manage storage pools for a vSnap server. You can also manage the cache and the log files for the server.

Managing disks

The vSnap server creates a storage pool by using the disks that are provisioned to the vSnap server. In the case of virtual deployments, the disks can be RDM or virtual disks provisioned from datastores on any backing storage. In the case of physical deployments, the disks can be local or be attached to the physical server in a storage area network (SAN). The local disks might already have external redundancy enabled via a hardware Redundant Array of Independent Disks (RAID) controller, but if not, the vSnap server can create RAID-based storage pools for internal redundancy.



Attention: Disks that are attached to vSnap servers must be thick provisioned. If disks are thin provisioned, the amount of free space in the storage pool might not be adequately reported. This situation might lead to data corruption if the underlying datastore runs out of space.

After a disk is added to a storage pool, do not remove the disk. Removing the disk will corrupt the storage pool.

If the vSnap server was deployed as part of a virtual appliance, the appliance already contains a 100 GB starter virtual disk. For instructions about managing this disk, see the [Blueprints](#). You can add more disks before or after creating a pool and accordingly use them to create a larger pool or expand an existing pool. If job logs report that a vSnap server is reaching its storage capacity, additional disks can be added to the vSnap pool. Or you can create an SLA policy and specify that backup operations use an alternative vSnap server as the target.

You can prevent data corruption, which can occur when a VMware datastore on a vSnap server reaches its capacity. Create a stable environment for virtual vSnap servers that use RAID configurations and utilize thick provisioned VMDKs. By replicating data to external vSnap servers, you can provide additional protection.

A vSnap server will become invalidated if the vSnap pool is deleted or if a vSnap disk is deleted. All data on the vSnap server will be lost. If your vSnap server becomes invalidated, you must unregister the vSnap server by using the IBM Spectrum Protect Plus interface, and then run the maintenance job. When the maintenance job is complete, register the vSnap server again.

Enabling encryption

To enable encryption of backup data on a vSnap server, select **Initialize with encryption enabled** when you initialize the server. Encryption settings cannot be changed after the server is initialized and a pool is created. All disks of a vSnap pool use the same encryption key file, which is generated upon pool creation. Data is encrypted when at rest on the vSnap server.

vSnap encryption utilizes the following algorithm:

Cipher name

Advanced Encryption Standard (AES)

Cipher mode

xts-plain64

Key

256 bits

Linux Unified Key Setup (LUKS) header hashing

sha256

Managing encryption keys

The disk encryption key files that are generated during pool creation are stored under the directory `/etc/vsnap/keys/` on each vSnap server. For disaster recovery purposes, back up the key files manually to another location outside of the vSnap server. After a pool is created, use the following commands as the `serveradmin` user to copy the keys to a temporary location and then copy them to a secure backup location outside the vSnap host. Complete the following steps:

1. Create a directory to which the keys will be backed up:

```
$ mkdir /tmp/keybackup-$(hostname)
```

2. Copy the key files to the temporary location:

```
$ sudo cp -r /etc/vsnap/keys /tmp/keybackup-$(hostname)
```

3. Copy the `keybackup-<hostname>` directory to a secure backup location outside of the vSnap host.

Detecting disks

If you add disks to a vSnap server, use the command line or the IBM Spectrum Protect Plus user interface to detect the newly attached disks.

Command line: Run the `$ vsnap disk rescan` command.

User interface: In the navigation panel, click **System Configuration** > **Storage** > **vSnap servers**. Then, click **Manage**. Open the **Storage disk** tab, and then click **Rescan**.

Showing disks

To view a list of all disks in the vSnap system, run the `$ vsnap disk show` command.

The USED AS column in the output shows whether each disk is in use. Any disk that is unformatted and unpartitioned is marked as unused. All other disks are marked as used.

Only disks that are marked as unused can be used to create a storage pool or be added to a storage pool. If a disk that you plan to add to a storage pool is not marked as unused, it might be because the disk was previously in use and thus contains remnants of an older partition table or file system. You can correct this issue by using system commands like **parted** or **dd** to wipe the disk partition table.

Showing storage pool information

To view information about each storage pool, run the `$ vsnap pool show` command.

Creating a storage pool

If you completed the simple initialization procedure that is described in “[Completing a simple initialization](#)” on [page 53](#), a storage pool was created automatically and the information in this section is not applicable.

To complete an advanced initialization, use the `vsnap pool create` command to create a storage pool manually. Before you run the command, ensure that one or more unused disks are available as described

in [“Showing disks”](#) on page 63. For information about available options, use the **--help** option for any command or subcommand.

Specify a display name for the pool and a list of one or more disks. If no disks are specified, all available unused disks are used. You can enable compression and deduplication for the pool during creation. You can also update the compression and deduplication settings later by using the **vsnap pool update** command.

The pool type that you specify during the creation of the storage pool specifies the redundancy of the pool:

raid0

This is the default option when no pool type is specified. If this option is used, vSnap assumes that your disks have external redundancy. This setting is appropriate, for example, if you use virtual disks on a datastore that is backed by redundant storage. In this case, the storage pool has no internal redundancy.

After a disk is added to a raid0 pool, the disk cannot be removed. If you remove the disk, the pool becomes unavailable. This issue can be resolved only by destroying and re-creating the pool.

raid5

When you select this option, the pool is comprised of one or more RAID5 group, each consisting of three or more disks. The number of RAID5 groups and the number of disks in each group depend on the total number of disks that you specify during pool creation. Based on the number of available disks, vSnap uses values that maximize total capacity while also helping to optimize redundancy of metadata.

raid6

When you select this option, the pool is comprised of one or more RAID6 group, each consisting of four or more disks. The number of RAID6 groups and the number of disks in each group depend on the total number of disks that you specify during pool creation. Based on the number of available disks, vSnap uses values that maximize total capacity while also helping to optimize redundancy of metadata.

Expanding a storage pool

Before you expand a pool, ensure that one or more unused disks are available as described in [“Showing disks”](#) on page 63.

Use the command line or the IBM Spectrum Protect Plus user interface to expand a storage pool.

Command line: Run the **\$ vsnap pool expand** command. For information about available options, use the **--help** option for any command or subcommand.

User interface: In the navigation panel, click **System Configuration > Storage > vSnap servers**. Then, click **Manage**. Open the **Storage disk** tab. The tab displays all unused disks that are detected on the system. Select one or more disks and click **Save** to add them to the storage pool.

Managing the cache and log for storage pools

To store cache and log data for vSnap storage, use solid-state drive (SSD) flash or non-volatile memory express (NVMe) disks. By adding cache and log space to storage pools, you can help to optimize the performance of the vSnap server by decreasing redundant input and output (I/O) to the server. For more information about configuring cache and log space for storage pools, see the [IBM Spectrum Protect Plus Blueprints](#).

You must use the command line to add or remove the cache and log. Because the cache and log do not store data permanently, you can remove them when the pool is online. However, ensure that no backup, restore, or replication operations are occurring before you issue the remove command.

Use the following commands to add and remove the cache or log. For information about the available options for a command, use the **--help** option. For examples of these commands as used in vSnap installation and configuration steps, see the [IBM Spectrum Protect Plus Blueprints](#).



Attention: Do not remove the devices that are providing space for the log and cache from the vSnap system without first removing the log and cache from the storage pool by using the appropriate remove command.

- **vsnap pool addcache**
- **vsnap pool addlog**
- **vsnap pool removecache**
- **vsnap pool removelog**

Network management

Configure and administer network services for a vSnap server.

The network on a vSnap server can be modified through the command line interface (CLI) through use of the **network** command. Additional information can be obtained by using the **--help** option after any command.

Showing network interface information

Run the **show** command to list network interfaces and the services that are associated with each interface:

```
$ vsnap network show
```

By default, the following vSnap services are available on all network interfaces:

mgmt

This service is used for management traffic between IBM Spectrum Protect Plus and vSnap.

repl

This service is used for data traffic between vSnap servers during replication.

nfs

This service is used for data traffic when backing up data using NFS.

smb

This service is used for data traffic when backing up data using SMB/CIFS.

iscsi

This service is used for data traffic when backing up data using iSCSI.

Modifying services associated with network interfaces

Run the **update** command to modify services that are associated with an interface. For example, if you are using a dedicated interface for data traffic to improve performance.

```
$ vsnap network update
```

The following options are required:

--id <id>

Enter the ID of the interface to update.

--services <services>

Specify all or a comma-separated list of services to enable on the interface. The following are valid values: **mgmt**, **repl**, **nfs**, **smb**, and **iscsi**.

If a service is available on more than one interface, IBM Spectrum Protect Plus can use any one of the interfaces.

Ensure that the **mgmt** service remains enabled on the interface that was used to register the vSnap server in IBM Spectrum Protect Plus.

Certificate management

You can manage your unique self-signed vSnap certificate in the IBM Spectrum Protect Plus environment.

Managing vSnap certificates

Beginning with IBM Spectrum Protect Plus version 10.1.11, each vSnap generates a unique self-signed certificate during the initial registration or deployment of the vSnap server. The certificate is configured with a hostname that is automatically detected during the initialization.

- The following hostname are embedded in the certificate by default:

Common Name (CN)

This is set to the fully qualified domain name (FQDN) of the vSnap server. Determine the **Common name** by using the following command:

```
hostname --fqdn
```

Subject Alternative Names (SAN)

Determine the **Short name** and **IP address** by using the following commands:

Note: When registering a vSnap in IBM Spectrum Protect Plus server, the vSnap certificate must be pasted or uploaded. The hostname or IP of the vSnap as entered in the IBM Spectrum Protect Plus UI must exactly match one of the SANs embedded in the vSnap certificate.

```
$ hostname
```

```
$ hostname -I
```

- Refer to the inline help on the vSnap server using the following commands:

```
$ vsnap system cert show --help
```

```
$ vsnap system cert regenerate --help
```

- To view the current certificate in PEM format, use the following command:

```
$ vsnap system cert show
```

This can be used to obtain the certificate that should be pasted or uploaded in the IBM Spectrum Protect Plus UI while registering a vSnap.

- If the existing CN or SAN in the certificate are incorrect, use the following command to regenerate a new self-signed certificate with the correct names.

```
$ vsnap system cert regenerate --hostnames <list_of_comma_separated_hostnames> --ipaddr  
<optional_list_of_comma_separated_IPs>
```

For example:

```
vsnap system cert regenerate --hostnames "vsnap1.example.com,vsnap1" --ipaddr "10.11.128.1"
```

- Alternatively, if you want to use a custom CA-signed certificate, obtain the necessary certificate and key files (in PEM format) and place them at the following locations:
 - The certificate (.crt file) must be placed under `/etc/vsnap/ssl/spp-vsnap.crt`
 - The private key (.key file) must be placed under `/etc/vsnap/ssl/spp-vsnap.key`
- After regenerating or replacing the certificate, the vSnap API service must be restarted by using the following command:

```
$ sudo systemctl restart vsnap-api
```

- Check if the new certificate is installed correctly by using the following command:

```
$ vsnap system cert show
```

Installing kernel headers and tools

Kernel headers and tools are not installed by default. If you plan to compile and use custom drivers, modules, or other software, install the appropriate kernel header or tool on the vSnap server.

About this task

When a vSnap server is installed or updated on RHEL 8 or later, a compatible Linux kernel version 4.18 is used. Kernel headers and tools headers and tools associated with the kernel are not installed. If you plan to compile or use custom drivers, modules, or other software, you must install the kernel packages. The Red Hat Package Manager (RPM) installers for the kernel headers and tools are available in the vSnap installation directory.

Procedure

1. Log on to the vSnap server as the `serveradmin` user. The initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus”](#) on page 96.
2. To determine the Linux kernel version, open a command line and issue the following command:

```
$ uname -r
```

The output is displayed, where `xxxx` represents the revision number of the kernel:

```
$ 3.10.xxxx
```

3. Navigate to this directory:

```
$ cd /opt/vsnap/config/pkgs/kernel/
```

4. In the directory, locate the `xxxxxxxx.rpm` file, which is the package to be installed. Be sure that the correct package is identified for the installed Linux kernel version. To install the kernel header or tool, issue the following command:

```
$ sudo yum localinstall xxxxxxxx.rpm
```

Results

The kernel header or tool is installed.

Troubleshooting vSnap servers

The vSnap servers in an IBM Spectrum Protect Plus environment provide disk storage for protecting data through backup and replication processes. The vSnap server configured in your environment might be used as the target, the source, or both server and target. In order to repair or replace a vSnap server that has failed, there are steps to follow so that the affected vSnap server is brought to a working state first so that backup and replication services can resume. This is to ensure minimum loss of data.

Preventing job failures by synchronizing vSnap and CIFS passwords

Communications between a vSnap server and a Common Internet File System (CIFS) share can be disrupted if credentials are shared, but passwords are out of sync. To prevent jobs from failing, you must synchronize the vSnap and CIFS passwords.

About this task

If you have already changed the vSnap system password, run the following command to synchronize the SMB/CIFS password with the system password. When prompted to enter a username, specify the account (for example, `serveradmin`) that is used to register the vSnap server in IBM Spectrum Protect Plus. When prompted to enter the password, specify the current system password for that account.

```
vsnap user resyncsmbpass
```

For information about how to synchronize passwords, see [“User management” on page 60](#).

Why is the vSnap server still offline?

After you restart the vSnap server, it continues to show a status of offline on the IBM Spectrum Protect Plus user interface.

If data deduplication is enabled or was previously enabled on a vSnap server, the deduplication table (DDT) is preloaded into memory during the vSnap server startup process. The DDT preloading process can introduce a 15-minute delay in the startup of the vSnap server services. During this time, the vSnap server shows with a status of `Offline` is displayed. Wait for at least 15 minutes for the process to be completed and for the vSnap server to return to the `Online` status. You can run the `vsnap_status` command to monitor the vSnap server services.

If any of the vSnap services is in the `activating` state, it means that the vSnap services are starting. When all services are in the `active` state, the vSnap server is back online.

How does SAN work with IBM Spectrum Protect Plus and a vSnap server?

VMware production or clone restore operations can use VMware SAN transport mode, which transports data in a storage area network (SAN) environment. To run a SAN-based restore operation, you can use the advanced setting **Enable Streaming (VADP) restore**, which was introduced in IBM Spectrum Protect Plus 10.1.5. This restore operation option is set by default. Coupled with this option, you can specify SAN transport mode in the VADP proxy options for a particular site.

By using the SAN transport mode, you can restore your data by using SAN transport for the VADP transport method to read/write to the datastore over the SAN. The logical unit numbers (LUNs) that comprise that datastore must be mapped to the machine by running an initial backup. This backup operation uses the zone and LUN mask as if they were members of the vSphere cluster to access the datastore over the SAN.

Tip: To view the advanced options when you are running a production or clone restore operation, switch the job options from **Default Setup** to **Advanced Setup**.

IBM Spectrum Protect Plus restores data by creating a datastore that vSphere detects, then a storage vMotion back to the target datastore is initiated. IBM Spectrum Protect Plus does not restore data by writing directly to the datastore. For this reason, using the SAN transport mode as a communication method for block-level incremental forever processing has fewer benefits. However, for initial full backup operations, by using SAN as a transport method, works well.

For information about how to set up and run a VMware restore job, see [“Restoring VMware data” on page 232](#).

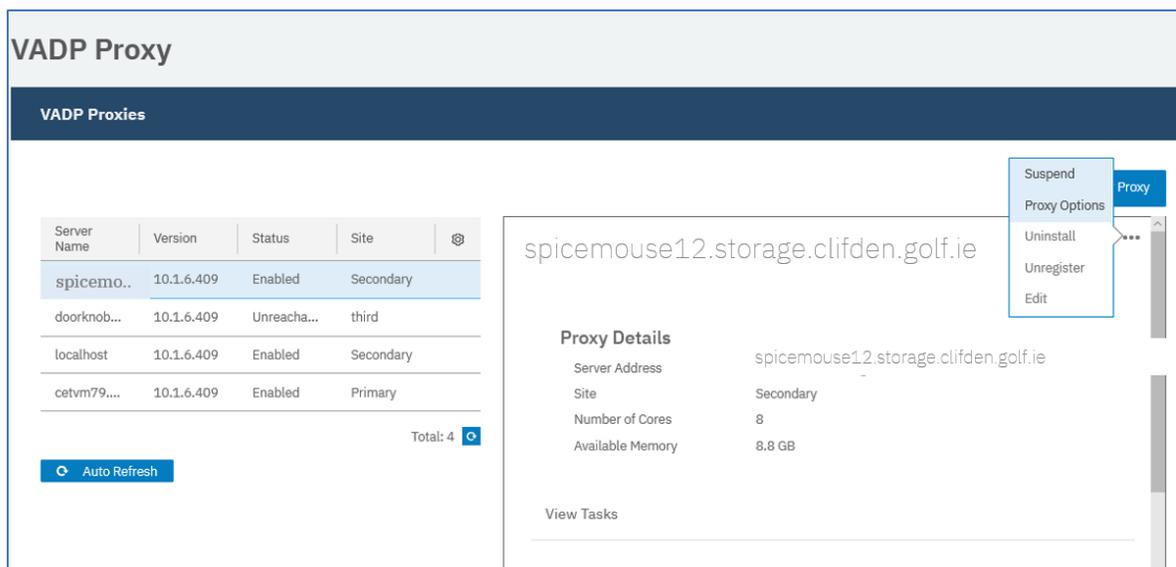
Communication

In IBM Spectrum Protect Plus, SAN backup is available through a physical proxy. Data transfer from storage to proxy is through the SAN. Communication from the proxy to the vSnap server is through the Network File System (NFS) protocol. The proxy and vSnap server can be installed on the same physical or virtual server. Review the proxy and vSnap server system requirements.

Specifying SAN as a data transport mode

To specify SAN as a transport mode, follow these steps:

1. Go to **System Configuration > VADP Proxy**. The **VADP Proxy** page opens.
2. From the table, select the server whose settings you want to edit. The **Proxy Details** pane shows the details for that server.
3. Click the actions icon  and select **Proxy Options**. The **Set VADP Proxy Options** dialog opens.



Server Name	Version	Status	Site
spicemo..	10.1.6.409	Enabled	Secondary
doorknob...	10.1.6.409	Unreacha...	third
localhost	10.1.6.409	Enabled	Secondary
cetvm79....	10.1.6.409	Enabled	Primary

Total: 4

Auto Refresh

Proxy Details

Server Address: spicemouse12.storage.clifden.golf.ie

Site: Secondary

Number of Cores: 8

Available Memory: 8.8 GB

View Tasks

4. From the **Transport Modes** list, select SAN.

Tip: When selection options include multiple transport modes, the first listed mode will be used. If that mode cannot be used, the next transport mode listed for that selection option will be used for transporting the data.

5. Click **Save**.

How do I repair a failed source vSnap in an IBM Spectrum Protect Plus environment?

The vSnap servers in an IBM Spectrum Protect Plus environment provide disk storage for protecting data through backup and replication processes. You can repair and replace a failed vSnap server that is configured in your IBM Spectrum Protect Plus environment to act as the *source* for backup and replication services. The source vSnap server must be repaired so that backup and replication services can resume.

Before you begin

Important: It is assumed that all vSnap servers in the environment are protected by replication. If a vSnap server is not replicated and it fails, it cannot be recovered to a state that would allow it to continue as a disk storage source or target. In the absence of replication processes, you must create a new vSnap server and set up service level agreement (SLA) policies. When you run the policies, a new full backup process runs to the new vSnap server.

To determine which type of repair process is applicable to your vSnap server, see [technote 1103847](#).

About this task

Important: Do not unregister or delete the failed vSnap server from IBM Spectrum Protect Plus. The failed vSnap server must remain registered for the replacement procedure to work correctly.

This procedure establishes a new source vSnap server in your IBM Spectrum Protect Plus environment to replace the failed source vSnap server. The new source vSnap server will contain only the most recent recovery points.

Note: The version of the new vSnap server must match the version of the deployed IBM Spectrum Protect Plus appliance.

Procedure

1. Log in to the target vSnap server console with the ID `serveradmin` by using Secure Shell (SSH) protocol.

Enter the following command: `$ ssh serveradmin@MGMT_ADDRESS`

For example, `$ ssh serveradmin@10.10.10.2`

2. Obtain the ID of the failed source vSnap server by opening a command prompt and entering the following command:

```
$ vsnap partner show
```

The output is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
MGMT ADDRESS: 10.10.10.1
API PORT: 8900
SSH PORT: 22
```

3. Verify that the MGMT ADDRESS is the address of the failed source vSnap server. Take note of the failed source vSnap server's ID number.
4. In the environment with the source vSnap server, install a new vSnap server of the same type and version, and with the same storage allocation, as the failed source vSnap server.

For instructions about installing a vSnap server, see [Installing a physical vSnap server](#).

Important: Do not register the new vSnap server with IBM Spectrum Protect Plus. Do not use the **Add vSnap server** wizard.

- a) You will first need to initialize the vSnap server with the following command:

```
$ vsnap system init --skip_pool --id partner_id
```

For example: `$ vsnap system init --skip_pool --id 12345678901234567890123456789012` using the failed source vSnap partner ID. A message indicates when the initialization is completed.

Note: This command is different to the vSnap initialization command listed in the IBM Documentation and in the Blueprints.

5. Complete the vSnap server and pool creation process as outlined in *Chapter 6: vSnap Server Installation and Setup* in the [Blueprints](#).
6. Place the new source vSnap server into maintenance mode by entering the following command:

```
$ vsnap system maintenance begin
```

Placing the vSnap server into maintenance mode suspends operations such as snapshot creation, data restore jobs, and replication operations.

7. Initialize the new source vSnap server with the failed source vSnap server's partner ID. Enter the following command:

```
$ vsnap system init --id partner_id
```

The following command is an example: `$ vsnap system init --id 12345678901234567890123456789012`

8. On the new source vSnap server, add the partner vSnap servers. Each partner must be added separately. To add a partner, enter the following command:

```
$ vsnap partner add --remote_addr remote_ip_address --local_addr local_ip_address
```

where, *remote_ip_address* specifies the IP address of the source vSnap server, and *local_ip_address* specifies the IP address of the new source vSnap server.

The following command is an example:

```
$ vsnap partner add --remote_addr 10.10.10.2 --local_addr 10.10.10.1
```

9. When prompted, enter the user ID and password for the target vSnap server. Informational messages indicate when the partners are created and updated successfully.
10. Create a repair task on the new source vSnap server by entering the following command:

```
$ vsnap repair create --async
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: N/A
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: PENDING
MESSAGE: The repair has been scheduled
```

11. Monitor the number of volumes that are involved in the repair operation by entering the following command:

```
$ vsnap repair show
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: 3
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: ACTIVE
MESSAGE: Created 0 volumes. There are 3 primary volumes that have recoverable snapshots,
the latest snapshot of each will be restored. Restoring 3 snapshots: 3 active, 0 pending, 0
completed, and 0 failed
```

The number of volumes that are involved in the repair operation is indicated in the TOTAL VOLUMES field.

12. Monitor the status of the repair task by viewing the repair.log file on the new source vSnap server, in the following directory `/opt/vsnap/log/repair.log`. Alternatively, you can enter the following command:

```
$ vsnap repair show
```

The output of this command is similar to the previous example. The following status messages can be displayed during the repair process:

- STATUS: PENDING indicates that the repair job is about to run.
- STATUS: ACTIVE indicates that the repair job is active.

- STATUS: COMPLETED indicates that the repair job is completed.
 - STATUS: FAILED indicates that the repair job failed and must be resubmitted.
13. During the repair operation, run the vsnap repair show command to verify when the status is COMPLETED.

```
$ vsnap repair session show
```

The output of this command is similar to the following example:

```
ID: 1 RELATIONSHIP: 72b19f6a9116a46aae6c642566906b31
PARTNER TYPE: vsnap
LOCAL SNAP: 1313
REMOTE SNAP: 311
STATUS: ACTIVE
SENT: 102.15GB
STARTED: 2019-11-01 15:51:18 UTC
ENDED: N/A
Created 0 volumes.
There are 3 replica volumes whose snapshots will be restored on next replication.
```

A session for each volume involved in the repair operation is displayed.

Periodically issue the `$ vsnap repair session show` command to ensure that the amount of data being sent for each volume is increasing in increments. As the sessions finish you will see the status change to COMPLETED. When all the sessions finish, issue the `$ vsnap repair session show` command to verify that the overall status is COMPLETED. A final message indicating the number of volumes for which snapshots were restored is displayed. The message output is similar to the following example:

```
Created 0 volumes.
There are 3 primary volumes that have recoverable snapshots, the latest snapshot of each
will be restored.
Restored 3 snapshots.
```

14. For any snapshots that are not restored and that indicate a FAILED status, resubmit the repair process by entering the following command:

```
$ vsnap repair create --async --retry
```

15. When the repair process reports a COMPLETED status, you can resume normal operations for the vSnap server by moving it out of maintenance mode. To resume normal processing, enter the following command:

```
$ vsnap system maintenance complete
```

16. Remove saved SSH host keys from the repaired source vSnap server and the target vSnap servers.

Run the following commands on both the source and target vSnap servers:

```
$ sudo rm -f /home/vsnap/.ssh/known_hosts
```

```
$ sudo rm -f /root/.ssh/known_hosts
```

Removing the SSH keys ensures that subsequent replication transfers do not produce errors that result from the changed host key of the repaired vSnap server.

17. Restart the vSnap service on the replaced server by entering the following command:

```
$ sudo systemctl restart vsnap
```

18. Click **System Configuration > Storage > vSnap servers** to verify that the new vSnap server is correctly registered, as follows:

- If the new vSnap server is using the same host name or IP address for registration, no change is required.
- If the new vSnap server is using a different host name or IP address for registration, click **Edit** to update the registration information.

19. To remove recovery points that are no longer available on the source vSnap server, start a maintenance job from the IBM Spectrum Protect Plus user interface.

For instructions, see [“Creating jobs and job schedules”](#) on page 432.

Tip: You might see informational messages that are similar to the following example:

```
CTGGA1843 storage snapshot spp_1004_2102_2_16de41fcbc3 not found on live Storage2101
Snapshot Type vsnap
```

20. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job. For instructions, see [“Creating jobs and job schedules”](#) on page 432.

Results

The source vSnap server has been repaired with only the most recent recovery points. The next backup job that runs as part of an SLA will back up data incrementally. If you create a restore job, only the most recent recovery point will be available in the backup repository. All other recovery points will be available in the replication repositories, and in the object storage and archive storage repositories if applicable to your environment.

How do I repair a failed target vSnap in an IBM Spectrum Protect Plus environment?

The vSnap servers in an IBM Spectrum Protect Plus environment provide disk storage for protecting data through backup and replication processes. You can repair and replace a failed vSnap server that is configured in your IBM Spectrum Protect Plus environment to act as the *target* for backup and replication services. The source vSnap server must be repaired so that backup and replication services can resume.

Before you begin

Important: It is assumed that all vSnap servers in the environment are protected by replication. If a vSnap server is not replicated and it fails, it cannot be recovered to a state that would allow it to continue as a disk storage source or target. In the absence of replication processes, you must create a new vSnap server and set up service level agreement (SLA) policies. When you run the policies, a new full backup process runs to the new vSnap server.

About this task

Important: Do not unregister or delete the failed vSnap server from IBM Spectrum Protect Plus. The failed vSnap server must remain registered for the replacement procedure to work correctly.

This procedure establishes a new target vSnap server in your IBM Spectrum Protect Plus environment to replace the failed target vSnap server. The new target vSnap server will not contain any data but will be populated with the most recent recovery points during the next scheduled replication operation.

Requirement: The version of the new vSnap server must match the version of the deployed IBM Spectrum Protect Plus appliance.

To determine which type of repair process is applicable to your vSnap server, see [technote 1103847](#).

Procedure

1. Log in to the functioning vSnap server console with the ID `serveradmin` by using Secure Shell (SSH) protocol.

Enter the following command: `$ ssh serveradmin@MGMT_ADDRESS`

For example, `$ ssh serveradmin@10.10.10.1`

2. Obtain the ID of the failed vSnap server by opening a command prompt and entering the following command:

```
$ vsnap partner show
```

The output is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
MGMT ADDRESS: 10.10.10.2
API PORT: 8900
SSH PORT: 22
```

3. Verify that the MGMT ADDRESS is the address of the failed vSnap server. Take note of the failed vSnap server's ID number.
4. In the environment with the target vSnap server, install a new vSnap server of the same type and version, and with the same storage allocation, as the failed target vSnap server.

For instructions about installing a vSnap server, see [Installing a physical vSnap server](#).

Important: Do not register the new vSnap server with IBM Spectrum Protect Plus. Do not use the **Add vSnap server** wizard.

- a) Initialize the vSnap server with the following command:

```
$ vsnap system init --skip_pool --id <partner_id>
```

For example, to use the partner ID of the failed source vSnap server, issue the following command:

```
$ vsnap system init --skip_pool --id 12345678901234567890123456789012
```

A message indicates when the initialization is completed.

Tip: This command is different from the vSnap initialization command listed in the Blueprints.

5. Complete the vSnap server and pool creation process as outlined in *Chapter 6: vSnap Server Installation and Setup* in the [Blueprints](#).
6. Place the new vSnap server into maintenance mode by entering the following command:

```
$ vsnap system maintenance begin
```

Placing the vSnap server into maintenance mode suspends operations such as snapshot creation, data restore jobs, and replication operations.

7. Initialize the new target vSnap server with the failed target vSnap server's partner ID. Enter the following command:

```
$ vsnap system init --id <partner_id>
```

The following command is an example:

```
$ vsnap system init --id 12345678901234567890123456789012
```

8. On the new target vSnap server, add the partner vSnap servers. Each partner must be added separately. To add a partner, enter the following command:

```
$ vsnap partner add --remote_addr <remote_ip_address> --local_addr <local_ip_address>
```

where, *<remote_ip_address>* specifies the IP address of the source vSnap server, and *<local_ip_address>* specifies the IP address of the new target vSnap server.

The following command is an example:

```
$ vsnap partner add --remote_addr 10.10.10.1 --local_addr 10.10.10.2
```

9. When prompted, enter the user ID and password for the source vSnap server.
Informational messages indicate when the partners are created and updated successfully.
10. Create a repair task on the new source vSnap server by entering the following command:

```
$ vsnap repair create --async
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: N/A
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: PENDING
MESSAGE: The repair has been scheduled
```

11. Monitor the number of volumes that are involved in the repair operation by entering the following command:

```
$ vsnap repair show
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: 3
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: ACTIVE
MESSAGE: Creating 3 volumes for partner 670d61a10f78456bb895b87c45e20999
```

The number of volumes that are involved in the repair operation is indicated in the TOTAL VOLUMES field.

12. Monitor the status of the repair task by viewing the repair.log file on the new source vSnap server, in the following directory /opt/vsnap/log/repair.log. Alternatively, you can enter the following command:

```
$ vsnap repair show
```

The output of this command is similar to the previous example. The following status messages can be displayed during the repair process:

- STATUS: PENDING indicates that the repair job is about to run.
- STATUS: ACTIVE indicates that the repair job is active.
- STATUS: COMPLETED indicates that the repair job is completed.
- STATUS: FAILED indicates that the repair job failed and must be resubmitted.

13. During the repair operation, run the vsnap repair show command to verify when the status is COMPLETED.

```
$ vsnap repair session show
```

The final message indicates the number of volumes whose snapshots will be restored on the next replication, as follows:

```
Created 0 volumes.
There are 3 replica volumes whose snapshots will be restored on next replication.
```

14. For any snapshots that are not restored and indicate a FAILED status, resubmit the repair process by entering the following command:

```
$ vsnap repair create --async --retry
```

15. When the repair process reports a COMPLETED status, you can resume normal operations for the vSnap server by moving it out of maintenance mode. To resume normal processing, enter the following command:

```
$ vsnap system maintenance complete
```

16. Remove saved SSH host keys from the repaired source vSnap server and the target vSnap servers. Run the following commands on both the source and target vSnap servers:

```
$ sudo rm -f /home/vsnap/.ssh/<known_hosts>
```

```
$ sudo rm -f /root/.ssh/<known_hosts>
```

Removing the SSH keys ensures that subsequent replication transfers do not produce errors that result from the changed host key of the repaired vSnap server.

17. Restart the vSnap service on the replaced server by entering the following command.

```
$ sudo systemctl restart vsnap
```

18. Click **System Configuration > Storage > vSnap servers** to verify that the new vSnap server is correctly registered, as follows:

- If the new vSnap server is using the existing hostname or IP address for registration, no change is required.
- If the new vSnap server is using a different hostname or IP address for registration, click **Edit** to update the registration information.

19. To remove recovery points that are no longer available on the source vSnap server, start a maintenance job from the IBM Spectrum Protect Plus user interface.

Tip: You might see informational messages that are similar to the following example:

```
CTGGA1843 storage snapshot spp_1004_2102_2_16de41fcbc3 not found on live Storage2101  
Snapshot Type vsnap
```

20. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job.

Results

The target vSnap server has been repaired. A new backup job must be run on the source vSnap server before any additional action is taken on the new target vSnap server.

If a replication job is attempted on the new target vSnap server, a message is displayed as follows:

```
CTGGA0289 - Skipping volume <volume_id> because there are no new snapshots since last backup
```

After a new backup job is run on the source vSnap server, the next scheduled replication job replicates the recovery points that are created by the backup job. At this point, if you create a restore job, only the most recent recovery point will be available in the replication repository. If the target vSnap server was also acting as a copy source to object or archive storage, the replication job must first run on the target vSnap server before any additional copy operations can complete successfully. The first copy of data to object storage will be a full copy.

How do I repair a failed dual-role vSnap in an IBM Spectrum Protect Plus environment?

You can repair and replace a failed vSnap server that is configured in your IBM Spectrum Protect Plus environment to act as both the *source* and *target* for backup and replication services.

About this task

Important: Do not unregister or delete the failed vSnap server from IBM Spectrum Protect Plus. The failed vSnap server must remain registered for the replacement procedure to work correctly.

This procedure establishes a new vSnap server in your IBM Spectrum Protect Plus environment to replace the failed vSnap server. After the repair process is completed, the new vSnap server is recovered to a point where backup jobs can continue to back up incremental changes (no full backup required) and replication jobs can continue.

To determine which type of repair process is applicable to your vSnap server, see [technote 1103847](#).

Note: The version of the new vSnap server must match the version of the deployed IBM Spectrum Protect Plus appliance.

Procedure

1. Log in to the functioning vSnap server in your environment console with the ID `serveradmin` by using Secure Shell (SSH) protocol.

Enter the following command: `$ ssh serveradmin@MGMT_ADDRESS`

For example, `$ ssh serveradmin@10.10.10.2`

2. Obtain the ID of the failed vSnap server by opening a command prompt and entering the following command:

```
$ vsnap partner show
```

The output is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
MGMT ADDRESS: 10.10.10.1
API PORT: 8900
SSH PORT: 22
```

3. Verify that the MGMT ADDRESS is the address of the failed vSnap server. Take note of the failed vSnap server's ID number.
4. On the target vSnap server, install a new vSnap server of the same type and version, and with the same storage allocation, as the failed source vSnap server.

For instructions about installing a vSnap server, see [Installing a physical vSnap server](#).

Important: Do not register the new vSnap server with IBM Spectrum Protect Plus. Do not use the **Add vSnap server** wizard.

- a) Initialize the vSnap server with the following command:

```
$ vsnap system init --skip_pool --id partner_id
```

For example, to use the partner ID of the failed source vSnap server, issue the following command:

```
$ vsnap system init --skip_pool --id 12345678901234567890123456789012
```

A message indicates when the initialization is completed.

Tip: This command is different from the vSnap initialization command listed in the Blueprints.

5. Complete the vSnap server and pool creation process as outlined in *Chapter 6: vSnap Server Installation and Setup* in the [Blueprints](#).

6. Place the new vSnap server into maintenance mode by entering the following command:

```
$ vsnap system maintenance begin
```

Placing the vSnap server into maintenance mode suspends operations such as snapshot creation, data restore jobs, and replication operations.

7. Initialize the new target vSnap server with the failed target vSnap server's partner ID. Enter the following command to initialize the vSnap:

```
$ vsnap system init --id partner_id
```

The following command is an example: `$ vsnap system init --id 12345678901234567890123456789012`

8. On the new target vSnap server, add the partner vSnap servers. If there is more than one partner server, each partner must be added separately. To add a partner, enter the following command:

```
$ vsnap partner add --remote_addr remote_ip_address --local_addr local_ip_address
```

where, `remote_ip_address` specifies the IP address of the source vSnap server, and `local_ip_address` specifies the IP address of the new target vSnap server.

The following command is an example:

```
$ vsnap partner add --remote_addr 10.10.10.1 --local_addr 10.10.10.2
```

9. When prompted, enter the user ID and password for the source vSnap server.

Informational messages indicate when the partners are created and updated successfully.

10. Create a repair task on the new source vSnap server by entering the following command:

```
$ vsnap repair create --async
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: N/A
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: PENDING
MESSAGE: The repair has been scheduled
```

11. Monitor the number of volumes that are involved in the repair operation by entering the following command:

```
$ vsnap repair show
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: 6
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: ACTIVE
MESSAGE: Created 0 volumes
There are 3 replica volumes whose snapshots will be restored on next replication.
There are 3 primary volumes that have recoverable snapshots, the latest snapshot of each
will be restored.
The number of volumes that are involved in the repair operation are indicated in the TOTAL
VOLUMES field
```

12. Monitor the status of the repair task by viewing the repair.log file on the new source vSnap server, in the following directory /opt/vsnap/log/repair.log. Alternatively, you can enter the following command:

```
$ vsnap repair show
```

13. When the status of the repair operation is in the ACTIVE state, you can view the status of individual repair sessions by entering the following command:

```
$ vsnap repair session show
```

The output is similar to this example:

```
ID: 1
RELATIONSHIP: 72b19f6a9116a46aae6c642566906b31
PARTNER TYPE: vsnap
LOCAL SNAP: 1313
REMOTE SNAP: 311
STATUS: ACTIVE
SENT: 102.15GB
STARTED: 2019-11-01 15:51:18 UTC
ENDED: N/A
```

View a session for each of the source volumes in the repair operation. The amount of data that is sent for each volume shows increasing incremental values until the process completes. The final message indicates the number of volumes whose snapshots will be restored by the next replication operation, as shown in this example:

```
Created 0 volumes. There are 3 replica volumes whose snapshots will be restored on next replication.
```

14. For any snapshots that are not restored and indicate a FAILED status, resubmit the repair process by entering the following command:

```
$ vsnap repair create --async --retry
```

15. When the repair process reports a COMPLETED status, you can resume normal operations for the vSnap server by moving it out of maintenance mode. To resume normal processing, enter the following command:

```
$ vsnap system maintenance complete
```

16. Optional: To view the total volumes and number of snapshots that were restored during the repair operation, run the show command for the vSnap server.

The output includes the following information:

- Total volumes lists the total number of volumes that were inspected during the repair operation. This list includes the source volumes (primary volumes) where the latest recovery point backup was restored, and target volumes (replica volumes) that are repopulated during upcoming replication operations as scheduled in SLAs.
- SNAPSHOTS RESTORED lists the number of source volumes that were restored.

17. Remove saved SSH host keys from the repaired source vSnap server and the target vSnap servers.

Run the following commands on both the source and target vSnap servers:

```
$ sudo rm -f /home/vsnap/.ssh/known_hosts
```

```
$ sudo rm -f /root/.ssh/known_hosts
```

Removing the SSH keys ensures that subsequent replication transfers do not produce errors that result from the changed host key of the repaired vSnap server.

18. Restart the vSnap service on the replaced server by entering the following command:

```
$ sudo systemctl restart vsnap
```

19. Click **System Configuration > Storage > vSnap servers** to verify that the new vSnap server is correctly registered, as follows:

- If the new vSnap server is using the same hostname or IP address for registration, no change is required.
- If the new vSnap server is using a different host name or IP address for registration, click **Edit** to update the registration information.

20. To remove recovery points that are no longer available on the source vSnap server, start a maintenance job from the IBM Spectrum Protect Plus user interface.

For instructions, see [“Creating jobs and job schedules”](#) on page 432.

Tip: You might see informational messages that are similar to the following example:

```
CTGGA1843 storage snapshot spp_1005_2102_2_16de41fcbc3 not found on live Storage2101  
Snapshot Type vsnap
```

21. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job. For instructions, see [“Creating jobs and job schedules”](#) on page 432.

Results

For primary backup data that is stored on the repaired vSnap server, the latest recovery point for primary backup data is now available. Subsequent backups to the repaired vSnap server continue to send only incremental changes since the last backup. For replicated data stored on the repaired vSnap server, no replicated data is available immediately after the repair. Subsequent replication jobs from the partner vSnap server will repopulate any backups that are created on the partner vSnap server after the repair process was completed. If a replication job is attempted on the partner vSnap server before a backup is completed on the partner vSnap server, a warning message is displayed indicating that there are no new snapshots since the last backup:

```
CTGGA0289 - Skipping volume <volume_id> because there are no new snapshots since last backup
```

If the repaired vSnap server was acting as a copy source to object or archive storage, a backup job must first be run on the repaired vSnap server before any additional copy operations will be successful. The first copy of data to object storage will be a full copy.

Chapter 5. Updating IBM Spectrum Protect Plus components

You can update the IBM Spectrum Protect Plus components to get the latest features and enhancements. Software patches and updates are installed by using the IBM Spectrum Protect Plus user interface or command-line interface for these components.

Important: Ensure that all vSnap servers are upgraded before you begin the IBM Spectrum Protect Plus server upgrade to a newer version.

For information about available update files and how to obtain them from an IBM download site, see [technote 6827871](#).

Before you update IBM Spectrum Protect Plus components, review the hardware and software requirements for the components to confirm any changes that might have occurred from previous versions.

Review the following restrictions and tips:

- The update process through the IBM Spectrum Protect Plus user interface updates IBM Spectrum Protect Plus features and the underlying infrastructure components including the operating system and file system. Do not use another method to update these components.
- Do not update any of the underlying components for IBM Spectrum Protect Plus unless the component is provided in an IBM Spectrum Protect Plus update package. Infrastructure updates are managed by IBM update facilities. The IBM Spectrum Protect Plus user interface is the primary means for updating IBM Spectrum Protect Plus features and underlying infrastructure components including the operating system and file system.

Before you update components, it is important that you back up your IBM Spectrum Protect Plus environment as described in [“Backing up the IBM Spectrum Protect Plus catalog”](#) on page 427.

Updating IBM Spectrum Protect Plus in a virtual appliance environment

Update IBM Spectrum Protect Plus install as a virtual appliance by using the administrative console.

Before you begin

Important:

- Ensure that all vSnap servers are upgraded before you begin the IBM Spectrum Protect Plus server upgrade to a newer version.
- Before you begin the update procedure, there must be at least 4.2 GB of disk space available in the /tmp directory of the IBM Spectrum Protect Plus server.

After the IBM Spectrum Protect Plus virtual appliance is updated, it cannot roll back to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of the IBM Spectrum Protect Plus appliance before you update to a new version of IBM Spectrum Protect Plus. If you later want to roll back IBM Spectrum Protect Plus to an earlier version, you must have a virtual machine snapshot. After the upgrade is completed successfully, remove the virtual machine snapshot.



Attention: Do not create virtual machine snapshots of external vSnap servers. After vSnap servers are updated, the servers cannot roll back to a previous version.

Managing updates

An IBM Spectrum Protect Plus environment includes the IBM Spectrum Protect Plus server and optionally one or more vSnap servers and VADP proxies. To help ensure that IBM Spectrum Protect Plus operates

normally, all components in the environment must be at the same version level. Review the instructions to carefully plan and complete the update process.

Before you begin

Important: Ensure that all vSnap servers are upgraded before you begin the IBM Spectrum Protect Plus server upgrade to a newer version.

Complete the following steps:

1. Plan a maintenance and verification period for the update process. You can estimate the required time based on the number of components in the environment that must be updated.

The process of upgrading an IBM Spectrum Protect Plus environment depends on the number of components in the environment and network speeds of the locations involved. The following table contains the three IBM Spectrum Protect Plus components and the average time, in minutes, that it takes to apply the update and successfully restart the system.

Component	Time to update	Time to restart	Total
IBM Spectrum Protect Plus server	10	15	25
vSnap server	15	10 - 30	25 - 45
VADP proxy server	15	Not required.	15

2. Gather version information for the components in your environment and determine the version levels for the update process. Determine if the vSnap servers must be updated as part of the upgrade process.

3. Adjust the start times of scheduled inventory or maintenance jobs so that they will run after the maintenance and verification period is concluded.

4. End any restore or reuse jobs, including object storage restore jobs. If necessary, schedule these jobs after the maintenance and verification period is completed.

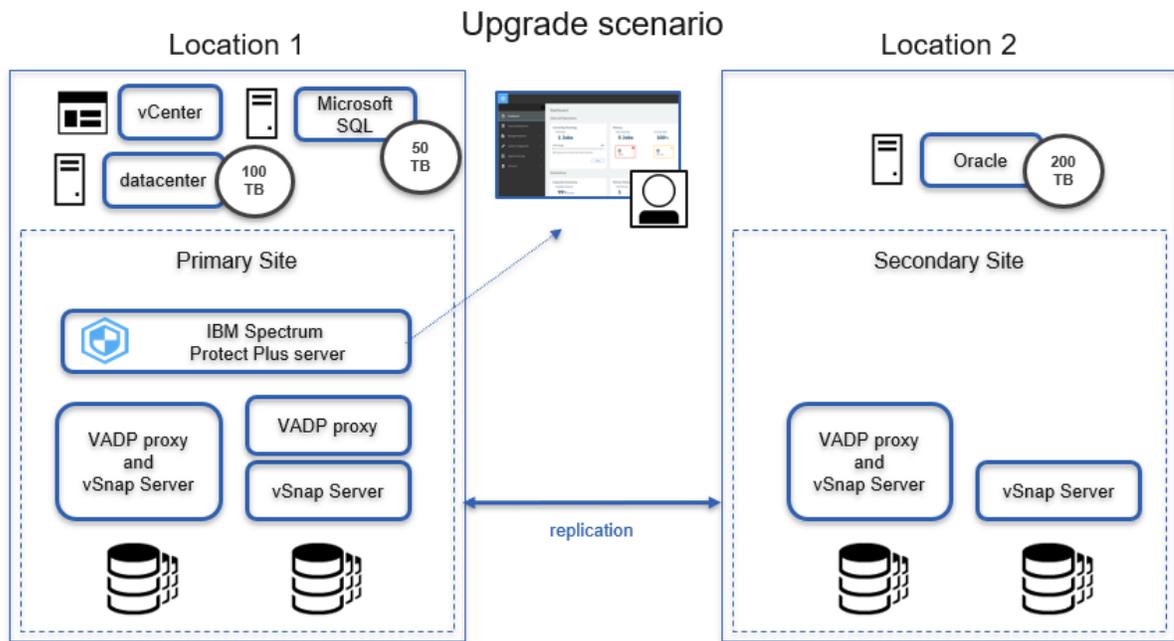
5. Pause any remaining jobs so that they do not run during the maintenance and verification period.

About this task

The procedure is based on an example environment, which includes the following components:

- 1 IBM Spectrum Protect Plus server
- 4 vSnap servers, with all four servers having replication relationships
- 2 VADP proxies co-installed with two of the vSnap servers
- 1 stand-alone VADP proxy

In the following figure, the components are displayed at their respective sites, Location 1 and Location 2:



Procedure

1. To prepare the system environment for the update process, complete the following steps:
 - a) In the navigation panel, click **Manage Protection > Policy Overview** and then click the **Add SLA Policy** button.
 - b) In the **New SLA Policy** pane, enter a policy name and click the radio button that includes the word **Catalog**. Click **Save**.
 - c) Select the **Disable Schedule** checkbox and specify an appropriate retention period. From the **Target Site** list, select the site that will contain the catalog backup.
 - d) Optionally, specify other options for the backup job. Click **Save**.
 - e) In the navigation panel, click **Manage Protection > IBM Spectrum Protect Plus > Backup**.
 - f) In the **SLA Policy** pane, select the policy that you created. Click **Save**.
 - g) The policy is displayed in the **SLA Policy Status** pane. If it does not appear automatically, click the refresh button.
 - h) To initiate the catalog backup, click **Actions** and then click **Start**.
 - i) Verify completion of the catalog backup job. In the navigation panel, click **Jobs and Operations** to verify that the catalog backup job was completed successfully.
 - j) Pause all scheduled jobs. In the navigation panel, click **Jobs and Operations** and click the **Schedule** tab. Click **Pause All Schedules**. The status for all scheduled jobs will change to **Held**.
 - k) Pause all scheduled jobs. In the navigation panel, click **Jobs and Operations** and click on the **Schedule** tab. Click **Pause All Schedules**. The status for all scheduled jobs changes to **Held**.
 - l) To verify that no jobs are running, click the **Running Jobs** tab. If jobs are running, allow the jobs to complete processing.
2. To prepare for updating vSnap servers, review the IBM Spectrum Protect Plus Blueprints at [Blueprints](#). Each vSnap server in your environment must be updated to the same IBM Spectrum Protect Plus version level. To update the vSnap servers, complete the following steps:
 - a) Follow the steps for updating the operating system for vSnap servers, as described in [“Updating the operating system for a virtual vSnap server”](#) on page 88.

Important: You must rename the downloaded ISO file as described in the procedure and move the file to the /tmp directory on the vSnap server if you wish to update the operating system.

- b) Complete the steps for updating a vSnap server, as described in [“Updating a vSnap server”](#) on page 89.

Tip: After you update a vSnap server, it can take 15 minutes longer than in previous versions to restart the vSnap server. For more information, see [Technote 3531159](#).

3. Update the IBM Spectrum Protect Plus server by completing the following steps:
 - a) Optional: If the IBM Spectrum Protect Plus server is deployed virtually, take a snapshot of the appliance in the appropriate hypervisor interface.
 - b) Update the IBM Spectrum Protect Plus server. Follow steps 1 through 6 in the [“Updating the IBM Spectrum Protect Plus server”](#) on page 85 topic. Do not release the schedule or any jobs that are held as indicated in the last two steps.
 - c) Log back in to the IBM Spectrum Protect Plus server.

4. Update VADP proxies. After you update the IBM Spectrum Protect Plus server, VADP proxies are updated automatically. However, the proxies might not be updated immediately. Depending on the number of proxies configured, it might take several minutes to an hour to update all proxies.

To update the VADP proxies immediately, follow the steps in the [“Updating VADP proxies”](#) on page 91 topic.

5. Verify that all components were updated successfully by completing the following steps:
 - a) Using the `serveradmin` account, log on to the IBM Spectrum Protect Plus administrative console. Follow the steps in [“Logging on to the administrative console”](#) on page 180.
 - b) Click **Product Information**. In the table, verify that the following items have the same version level: `spp-release`, `vsnap`, `vsnap-dist`, `vadp`, and `vadp-dist`.
 - c) Log out of the IBM Spectrum Protect Plus administrative console.
 - d) Load the IBM Spectrum Protect Plus splash screen by opening a supported browser and entering the following URL:

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

where *hostname* is the IP address of the virtual machine where the application is deployed.

- e) Verify that the version and build on the splash screen match the `spp-release` that was displayed in the **Product Information** section of the administrative console.
 - f) To verify that a maintenance job can be completed successfully in the updated environment, in the navigation panel, click **Jobs and Operations > Schedule**. Click the options icon  next to Maintenance Job and select **Start**. Monitor the job progress through the **Jobs and Operations** pane.
6. Release scheduled jobs and, optionally, remove the snapshot. Complete the following steps:
 - a) Release all schedules. In the navigation panel, click **Jobs and Operations > Schedule**. Click **Release All Schedules**.
 - b) Optional: If you took a snapshot of the IBM Spectrum Protect Plus virtual appliance, you can delete the snapshot of the IBM Spectrum Protect Plus server by using the hypervisor interface. Follow the instructions in the hypervisor documentation.

What to do next

If necessary, restart any jobs that were stopped or paused during the maintenance and verification period.

Updating the IBM Spectrum Protect Plus server

Use the IBM Spectrum Protect Plus administrative console to update the IBM Spectrum Protect Plus server in a virtual appliance environment. Updating IBM Spectrum Protect Plus can be run offline or online.

Before you begin

Important: Ensure that all vSnap servers are upgraded before you begin the IBM Spectrum Protect Plus server upgrade to a newer version.

You can update IBM Spectrum Protect Plus directly from two previous versions ($n-2$) to the current version (n). If you are using an older version, you must update at least to ($n-2$) version and then update to the current version.

Before you begin the update process, complete the following steps:

1. Ensure that your IBM Spectrum Protect Plus environment is backed up. For instructions, see [“Backing up the IBM Spectrum Protect Plus catalog”](#) on page 427.
2. For offline updates, download the IBM Spectrum Protect Plus update package that is named `<part_number>.iso` to a directory on the computer that is running the browser for the administrative console. The update file is installed first.
3. Ensure that no jobs are running before starting the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED. For instructions, see [“Pausing and resuming jobs”](#) on page 437.
4. Take a snapshot of the IBM Spectrum Protect Plus virtual appliance so that in the event of an update failure, the IBM Spectrum Protect Plus VM can be reverted.

For a list of download images, including the required operating system update for the virtual appliance, see [technote 6827871](#).

About this task

If you have access to the internet, you can choose to run the update procedure online. If you do not have internet access, you can run the update procedure offline.

Procedure

To update the IBM Spectrum Protect Plus virtual appliance, complete the following steps:

1. From a supported web browser, access the administrative console by entering the following address:

```
https://hostname:8090/
```

where *hostname* is the IP address of the virtual machine where the application is deployed.

2. In the login window, select one of the following authentication types in the **Authentication Type** list:

Authentication Type	Login information
IBM Spectrum Protect Plus	To log on as the IBM Spectrum Protect Plus superuser, enter the username and password. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.
System (recommended)	To log on as a system user, enter the <code>serveradmin</code> password. The default password is <code>sppDP758-SysXyz</code> . You are prompted to change this password during the first logon.

3. Click **Updates and Hotfix Management** to open the updates management page.

If you have access to the FTP site, public.dhe.ibm.com, the administrator console checks for available updates automatically and lists them.

4. Click **Run Update** to install the available updates.

- When the updates are installed successfully, go to Step 6.
- If you are planning to install an update from an ISO file, click **Click Here** to run the offline updates. Go to Step 5.

Note: If you want to run online updates but can see only the offline mode, check your internet connectivity and reattempt to access the FTP site, public.dhe.ibm.com.

5. Choose the update that you want to run, as follows:

- Online mode: Updates are listed automatically in the repository when they are made available. Click **Run Update**.
- Offline mode: Click **Choose file** to browse for the downloaded file. The file has an iso or rpm extension like this example, <filename>.iso. Click **Upload Update Image (or) Hotfix**. You can select only one update file at a time.

Important: There must be at least 4.2 GB of disk space available in the /tmp directory of the IBM Spectrum Protect Plus server.

If the update fails during the prerequisite check, you must either correct the error found and then reboot the IBM Spectrum Protect Plus appliance before re-attempting the update or revert the IBM Spectrum Protect Plus appliance, correct the prerequisite issue, and then re-attempt the update. When the update is completed, IBM Spectrum Protect Plus automatically restarts.

Important: After the IBM Spectrum Protect Plus update is completed, you must update the external vSnap and VADP proxy servers in your environment. For instructions see [“Updating vSnap servers” on page 88](#) and [“Updating VADP proxies” on page 91](#).

6. Clear the browser cache.

HTML content from previous versions of IBM Spectrum Protect Plus might be stored in the cache.

7. Start the updated version of IBM Spectrum Protect Plus.

8. In the navigation panel, click **Jobs and Operations**, and then click the **Schedule** tab.

Find the jobs that you paused.

9. Click the actions menu icon  for the job, and then click **Release Schedule**.

Related tasks

[“Updating vSnap servers” on page 88](#)

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Updating IBM Spectrum Protect Plus in a cloud environment

When IBM Spectrum Protect Plus is installed in a cloud environment, the steps to update the product components are similar to the steps for updating components in an on-premises environment.

About this task

The steps to update the components depend on the cloud environment that you are using: all on cloud or hybrid.

Hybrid environment

If you are updating IBM Spectrum Protect Plus in a hybrid cloud environment, the IBM Spectrum Protect Plus server is installed on-premises and the vSnap server is installed on the cloud.

To update the IBM Spectrum Protect Plus server on premises, follow the instructions in [“Updating the IBM Spectrum Protect Plus server” on page 85](#).

To update the vSnap server on the cloud, use a virtual private network (VPN) or other connection such as a Bastion host to access the host where the application is installed. To update the vSnap server, follow the instructions in [“Updating vSnap servers” on page 88](#).

All-on-cloud environment

If you are updating IBM Spectrum Protect Plus in an all-on-cloud environment, the IBM Spectrum Protect Plus server and the vSnap server are installed on the cloud.

To update these applications, use a VPN or other connection such as a Bastion host to access the host where the applications are installed, and then follow the instructions in [“Updating the IBM Spectrum Protect Plus server”](#) on page 85 and [“Updating vSnap servers”](#) on page 88.

Related concepts

[“IBM Spectrum Protect Plus on IBM Cloud”](#) on page 16

IBM Spectrum Protect Plus is available as a software offering in the IBM Cloud catalog or as an IBM Cloud for VMware Solutions service.

[“IBM Spectrum Protect Plus on the AWS cloud platform”](#) on page 17

IBM Spectrum Protect Plus on the Amazon Web Services (AWS) cloud platform is a data protection solution for users who want to protect databases that are running on AWS. In addition, users can protect virtual machines that are managed by VMware Cloud (VMC) on AWS while having the IBM Spectrum Protect Plus server installed on VMC and the vSnap server installed on an AWS Virtual Private Cloud (VPC).

[“IBM Spectrum Protect Plus on the Microsoft Azure cloud platform”](#) on page 18

IBM Spectrum Protect Plus on the Microsoft Azure cloud platform is a data protection solution for users who want to protect one or more databases that are running on Azure.

Updating IBM Spectrum Protect Plus by using the user interface

Use the IBM Spectrum Protect Plus user interface to update the product online.

Before you begin

Complete the following tasks before you start the update process:

- Log in to the IBM Spectrum Protect Plus as the superuser. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.
- Ensure that your IBM Spectrum Protect Plus environment is backed up before you run updates. For more information about backing up your environment, see [“Backing up the IBM Spectrum Protect Plus catalog”](#) on page 427.
- Ensure that no jobs are running during the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED.
- Clear the cache for the browser that you use to open the IBM Spectrum Protect Plus user interface. HTML content from previous versions of IBM Spectrum Protect Plus might be stored in the cache.

Procedure

To update IBM Spectrum Protect Plus, complete the following steps:

1. In the IBM Spectrum Protect Plus user interface, click the user menu  in the menu bar, and then click **Update IBM Spectrum Protect Plus**.
The installed versions are shown.
2. Click **Check for updates**.
3. Click the version that you want to update to, and then click **Proceed to update**.
4. Review the summary information, and then click **Update**.
5. Click **Yes** to confirm the update.

A window opens showing the update state for each IBM Spectrum Protect Plus component.

6. When the status of all components is **Running**, click **Reload** and then click **Reload** again in the confirmation dialog box.

The IBM Spectrum Protect Plus login page is opened with a message that the server is being brought up. When this message is no longer shown, you can log in to IBM Spectrum Protect Plus.

7. Log in to IBM Spectrum Protect Plus.
8. In the navigation panel, click **Jobs and Operations**, and then click the **Schedule** tab.
Find the jobs that you paused.
9. Click the actions menu icon  for the job, and then click **Release Schedule**.

Updating vSnap servers

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Before you begin

Important: Ensure that all vSnap servers are upgraded before you begin the IBM Spectrum Protect Plus server upgrade to a newer version.

You can update the IBM Spectrum Protect Plus and vSnap servers directly from two previous versions ($n-2$) to the current version (n). If you are using an older version, you must update at least to ($n-2$) version and then update to the current version.

Test restore jobs need to complete prior to initiating an update to vSnap. During a vSnap upgrade, a reboot will occur and any connected clients will experience a temporary disconnection. This disconnection may result in errors for any virtual machines or applications with active test mode restore. Additionally, jobs that are not completed or canceled when an update is initiated will not be visible once the update has completed. If jobs are not visible once the update has completed, re-run test restore jobs.

You might also be required to update the operating system for the vSnap servers prior to updating the servers. For operating system requirements, see [“Component requirements” on page 21](#).

To check the current version and operating system for your vSnap servers, complete the following steps:

1. Log on to the vSnap server as the `serveradmin` user. If you are using IBM Spectrum Protect Plus 10.1.1, log in by using the `root` account.
2. To check the vSnap server version and operating system, use the vSnap command-line interface to issue the following command:

```
$ vsnap system info
```

Ensure that no jobs that use the vSnap server are running during the update procedure. Pause the schedule for any jobs that do not have a status of IDLE or COMPLETED.

Updating the operating system for a physical vSnap server

If you have installed the vSnap server on a machine that is running Red Hat Enterprise Linux, you must update the operating system to version 7.5 or 7.6 before you update the vSnap server. For instructions about how to update the operating system, see the [Red Hat documentation](#).

Related tasks

[“Updating a vSnap server” on page 89](#)

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Updating the operating system for a virtual vSnap server

Updating the vSnap server operating system with the ISO file, provides you with the latest available patches and security updates. If the operating system is CentOS Linux version 7.4 or earlier, you must update the operating system before you update the vSnap server software. Updating the operating system is optional for version 7.5 or 7.6. An ISO file is downloaded and used to upgrade virtual vSnap servers.

Before you begin

You can update the IBM Spectrum Protect Plus and vSnap servers directly from two previous versions ($n-2$) to the current version (n). If you are using an older version, you must update at least to ($n-2$) version and then update to the current version.

Before you begin the update process, ensure that you have backed up your IBM Spectrum Protect Plus environment as described in “[Backing up the IBM Spectrum Protect Plus catalog](#)” on page 427. For information on obtaining the ISO file, see “[Updating the IBM Spectrum Protect Plus server](#)” on page 85.

Restriction: The ISO should not be used if updating a physical Red Hat Enterprise Linux server. It should only be used on OVA deployments.

Procedure

1. Download the ISO file `<part_number>.iso`. Move the ISO file to the `/tmp` directory on the vSnap server and rename the file to `spp_with_os.iso`.

```
$mv <part_number>.iso /tmp/spp_with_os.iso
```

Important: It is critical to rename the downloaded ISO file as described in this step and move it to the `/tmp` directory on the vSnap server if you wish to update the operating system.

2. Next, follow the instructions found in the “[Updating a vSnap server](#)” on page 89 topic. When the `<part_number>.run` file is executed, the installer will optionally update the operating system if `/tmp/spp_with_os.iso` is present.

Note: Verify that the correct installation file for Linux is downloaded. For upgrades, the image file that contains `vsnap-dist-el7-<version-build>` should be used for both CentOS and Red Hat Enterprise Linux 7.x. Download the associated **Passport Advantage Part Number** file that is named `<part_number>.run`.

One of the two following scenarios will occur depending on the presence of the ISO file.

- If the file is present, operating system packages are upgraded, then vSnap software is upgraded.
- If the file is not present, a message is displayed:

```
File /tmp/spp_with_os.iso is not present, skipping update of OS packages.  
To update OS packages, download the ISO file to /tmp/spp_with_os.iso and rerun this  
installer.
```

Then vSnap software is then is upgraded.

Once the installer completes, `/tmp/spp_with_os.iso` can be deleted.

Related tasks

“[Updating a vSnap server](#)” on page 89

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Updating a vSnap server

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Before you begin

You can update the IBM Spectrum Protect Plus and vSnap servers directly from two previous versions ($n-2$) to the current version (n). If you are using an older version, you must update at least to ($n-2$) version and then update to the current version.

Test restore jobs need to complete prior to initiating an update to vSnap. During a vSnap upgrade, a reboot will occur and any clients will experience a temporary disconnection. This disconnection may result in errors for any virtual machines or applications with active test mode restore. Additionally, jobs that are not completed or canceled when an update is initiated will not be visible once the update has completed. If jobs are not visible once the update has completed, re-run test restore jobs.

Before you begin the update process, complete the following steps:

1. Ensure that you have backed up your IBM Spectrum Protect Plus environment as described in “Backing up the IBM Spectrum Protect Plus catalog” on page 427.
2. Download the vSnap update file. The image name that contains *vsnap-dist-el8-
<version-build>* is for Red Hat Enterprise Linux 8.x and must be used for upgrades. The image name has an associated **Passport Advantage Part Number**. Download the corresponding run file, *<part_number>.run*, for the image name from Passport Advantage Online. For information about downloading files, see [technote 6827871](#).
3. Copy the upgrade file to a temporary location on the vSnap server.

Procedure

To update a vSnap server, complete the following steps:

1. Log on to the vSnap server as the `serveradmin` user.
2. From the directory where the *<part_number>.run* file is located, make the file executable by issuing the following command:

```
$ chmod +x <part_number>.run
```

3. Run the installer by issuing the following command:

```
$ sudo ./<part_number>.run
```

Alternatively, non-interactive installations or updates of vSnap may be initiated using the `noprompt` option. When this option is used, the vSnap installer will skip prompting for responses and assume an answer of "yes" to the following prompts:

- License agreement
- Kernel installation or update
- Reboot at the end of the installation or update if necessary

To use the `noprompt` option, issue the following command. Observe the deliberate space both before and after the double dashes:

```
$ sudo ./<part_number>.run -- noprompt
```

The vSnap packages are installed.

4. After the vSnap packages are installed, start the updated version of the vSnap server.
5. In the navigation panel, click **Jobs and Operations**, and then click the **Schedule** tab. Find the jobs that you paused.
6. From the **Actions** menu for the paused jobs, select **Release Schedule**.

Additional steps for updating virtual machines in Hyper-V Replica environments

Beginning with IBM Spectrum Protect Plus Version 10.1.5, you can protect virtual machines (VMs) that are enabled to use the Hyper-V Replica feature.

IBM Spectrum Protect Plus processes the data on the source and replicated instances of the VMs separately. For example, if a VM named VM1 is on the Hyper-V host named Host1 and the VM is replicated to Host2, IBM Spectrum Protect Plus assigns the IDs VM1@Host1 and VM1@Host2 to the VMs. You can then select one or both of the VMs for data protection.

Considerations for VMs that are defined in existing SLA policies

If you update IBM Spectrum Protect Plus, you might have to take additional steps to ensure that data protection continues for VMs that are currently included in your service level agreement (SLA) policies.

An SLA policy can *implicitly* or *explicitly* include a replicated VM. You might be required to update the SLA policy when you update to IBM Spectrum Protect Plus 10.1.5 or later.

An example of an SLA policy that implicitly includes a replicated VM is a scenario in which the policy protects all VMs on Host1, which contains the VM VM1. VM1 is replicated to Host2. In this scenario, a change to the SLA policy is not required after you update IBM Spectrum Protect Plus. The SLA policy creates a full backup of the instance of VM1 on Host2 and creates a new full backup of the instance of VM1 on Host1. Existing backups of VM1 on Host1 that were created before the update will expire based on the SLA policy retention settings.

An example of an SLA policy that explicitly includes a replicated VM is a scenario in which the policy protects VM1 on Host1, and VM1 is replicated to Host2. In this scenario, you must re-add the instance of the VM on each host to the SLA policy after you update IBM Spectrum Protect Plus.

Updating VADP proxies

Updating the IBM Spectrum Protect Plus virtual appliance automatically updates all the VADP proxies that are associated with the virtual appliance. In rare scenarios such as loss of network connectivity, you must update the VADP proxy manually.

Before you begin

Before you begin, ensure that you have backed up your IBM Spectrum Protect Plus environment as described in [“Backing up the IBM Spectrum Protect Plus catalog”](#) on page 427.

Note: Only VADP proxies registered with IBM Spectrum Protect Plus will be updated. If the VADP proxy is not registered with IBM Spectrum Protect Plus, the VADP component will not be updated.

Procedure

If a VADP proxy update is available for external proxies during a restart of the IBM Spectrum Protect Plus virtual appliance, the update will be automatically applied to any VADP proxy associated with an identity. To associate a VADP proxy with an identity, navigate to **System Configuration > VADP Proxy**. Click the ellipses icon ******* and select **Edit**. Select **Use existing user** and choose a previously entered identity in **Select user** for the VADP proxy server.

To update a VADP proxy manually, complete the following steps:

1. Navigate to the **System Configuration > VADP Proxy** page in IBM Spectrum Protect Plus.
2. The **VADP Proxy** page displays each proxy server. If a newer version of the VADP proxy software is available, an update icon  displays in the **Status** field.
3. Ensure that there are no active jobs that use the proxy, and then click the update icon .
The proxy server enters a suspended state and installs the latest update. When the update completes, the VADP proxy server automatically resumes and enters an enabled state.

When attempting to install a VADP proxy to a supported, stand-alone Linux deployment or when the account used to register the VADP proxy through IBM Spectrum Protect Plus is a non-root user, special instructions must be followed. Specifically, the username and password for the account used to register the VADP proxy must also exist on the machine to which the proxy is being installed or updated and have a matching `sudoers` configuration file. The `sudoers` configuration must allow the user to run commands without a password.

1. Log in to the VADP server as the `root` user.
2. In the case of a stand-alone VADP proxy deployed on a supported version of Linux, create a new user and assign a password. This is the account that will subsequently be used to register the VADP proxy

through the IBM Spectrum Protect Plus user interface. In this example, the variable *vadpuser* is the username used to register the VADP proxy.

```
# useradd vadpuser
# passwd vadpuser_password
```

In the case of updating a VADP proxy, verify that the user that was used to register the VADP server exists.

```
# cat /etc/passwd | grep vadpuser
```

3. When installing a stand-alone VADP proxy on a supported version of Linux, you may need to install the *nfs-utils* package if it is not already installed. Answer yes ("y") to the prompts.

```
# yum install nfs-utils
```

4. Next, create a *sudoers* configuration file in the */etc/sudoers.d/* directory. Write the "Defaults !requiretty" text to the file and save it by pressing CTRL+D on the keyboard when done.

```
# cat /etc/sudoers.d/vadpuser
Defaults !requiretty
vadpuser ALL=NOPASSWD: /tmp/cdm_guestapps_vadpuser/runcommand.sh
<<Press CTRL+D>>
```

5. Finally, set the appropriate permissions on the file.

```
# chmod 0440 vadpuser
```

What to do next

After you update the VADP proxies, complete the following action:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 220. The proxies are indicated in the job log by a log message similar to the following text: Run remote vmdkbackup of MicroService: http://<proxy <i>nodename</i> , IP: <i>proxy_IP_address</i>

Related tasks

[“Creating VADP proxies”](#) on page 227

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Related reference

[“Editing firewall ports”](#) on page 24

Use the provided examples as a reference for opening firewall ports on remote VADP proxy servers or application servers. You must restrict port traffic to only the required network or adapters.

Applying early availability updates

Early availability updates provide fixes for authorized program analysis reports (APARs) and minor issues between IBM Spectrum Protect Plus releases. These updates are available in bundles from the Fix Central Online website.

About this task

Early availability updates might not contain fixes for all IBM Spectrum Protect Plus components.

For instructions about how to obtain and install interim fixes, see the download information that is published when the fixes are available.

Chapter 6. Getting off to a quick start

To start using IBM Spectrum Protect Plus, you must define resources that you want to protect and create service level agreement (SLA) policies, also known as backup policies, for those resources. This getting started section provides these and other steps required to set up and start using IBM Spectrum Protect Plus to back up data. Other tasks such as copying and restoring data are discussed in detail in other areas of the documentation.

Ensure that you followed the instructions in the [IBM Spectrum Protect Plus Blueprints](#) to determine how to size, build, and place the components in your IBM Spectrum Protect Plus environment and that the tasks listed in the [“Deployment storyboard for IBM Spectrum Protect Plus”](#) on page 1 are complete.

As shown in the following table, the initial installation and configuration tasks are completed by the IBM Spectrum Protect Plus *infrastructure administrator*. By default, the `admin` user account is created for use by the infrastructure administrator to start the application for the first time.

Then, resource backup and restore tasks are completed by the *application administrator*. However, a single administrator might be responsible for all tasks in your environment.

Action	Owner	Description
Start IBM Spectrum Protect Plus	Infrastructure administrator and application administrator	<p>The infrastructure administrator starts the application for the first time by using the default <code>admin</code> user account with the password <code>password</code>. The administrator is prompted to reset the username and password for this account. The administrator cannot reset the user name to <code>admin</code>, <code>root</code>, or <code>test</code>.</p> <p>After the initial startup, the application administrator can start the application by using this user account, which is referred to as the IBM Spectrum Protect Plus superuser account, or another account that the infrastructure administrator creates.</p>
“Manage sites” on page 97	Infrastructure administrator	<p>A site is used to group vSnap servers based on a physical or logical location to help quickly identify and interact with backup data. A site is assigned to a vSnap server when the server is added to IBM Spectrum Protect Plus.</p> <p>The default sites are named Primary, Secondary, and Replication. You can also create a custom site.</p>

Action	Owner	Description
Create backup policies	Infrastructure administrator	<p>Backup policies define the parameters that are applied to backup jobs. These parameters include the frequency and retention of backups and the options to replicate data from one vSnap server to another and to copy backup data to secondary backup storage for longer-term protection.</p> <p>Backup policies also define the target site to for backing up data. A site can contain one or more vSnap servers.</p> <p>Backup policies are called SLA policies in IBM Spectrum Protect Plus.</p>
Create a user account for the application administrator	Infrastructure administrator	User accounts determine the resources and functions that are available to the user.
Add resources to protect	Application administrator	Resources are entities that you want to protect. After a resource is registered, an inventory of the resource is captured and added to the IBM Spectrum Protect Plus inventory.
Add resources to a job definition	Application administrator	Job definitions associate the resources that you want to protect with one or more SLA policies. The options and schedules that are defined in the SLA policies are used for backup jobs for the resources.
Start a backup job	Application administrator	Backup jobs are started as defined in the SLA policy that is associated with the job definition. You can also manually start a job.
Run a report	Application administrator	IBM Spectrum Protect Plus provides predefined reports that you can run with default parameters or modify to create custom reports.

Start IBM Spectrum Protect Plus

Start IBM Spectrum Protect Plus to begin using the application and its features.

Procedure

To start IBM Spectrum Protect Plus, complete the following steps:

1. In a supported web browser, enter the following URL:

```
https://hostname
```

The *hostname* value depends on whether IBM Spectrum Protect Plus installed as a set of OpenShift containers or as a virtual appliance.

Hostname for a container installation

The hostname must be in the following format:

```
instancename-spp.routerCanonicalHostname
```

where *instancename* is the name of the IBM Spectrum Protect Plus instance and *routerCanonicalHostname* is the external host name for the OpenShift router.

Hostname for a virtual appliance installation

The hostname is the IP address of the virtual machine where the application is deployed.

2. Enter your username and password to log on to IBM Spectrum Protect Plus.

If this is your first time logging on, the default username is `admin` and the password is `password`. You are prompted to reset the default username and password. You cannot reset the username to `admin`, `root`, or `test`.

This user account is the superuser account and is assigned the `SUPERUSER` role. This role is assigned to only one IBM Spectrum Protect Plus user. The `SUPERUSER` role provides the user with access to all IBM Spectrum Protect Plus functions. For more information about the superuser account, see [“Managing the superuser account”](#) on page 468.

3. Click **Sign In**.

4. If IBM Spectrum Protect Plus is installed on a virtual appliance and you are logging in for the first time, you are prompted to change the `serveradmin` password. The initial password is `sppDP758-SysXyz`. The `serveradmin` user is used to access the administrative console and the IBM Spectrum Protect Plus virtual appliance. The password for `serveradmin` must be changed before accessing the administrative console and IBM Spectrum Protect Plus virtual appliance.

The following rules are enforced when creating a new password:

- The minimum acceptable password length is 15 characters.
- There must be eight characters in the new password that are not present in the previous password.
- The new password must contain at least one character from each of the classes (numbers, uppercase letters, lowercase letters, and other).
- The maximum number of identical consecutive characters that are allowed in the new password is three characters.
- The maximum number of identical consecutive class of characters that are allowed in the new password is four characters.

Manage sites

A site is used to group the storage servers based on a physical or logical location to help quickly identify and interact with backup data. A site is assigned to a vSnap server when the server is added to IBM Spectrum Protect Plus.

About this task

Review the available sites by clicking **System Configuration > Storage > Sites** in the navigation panel and decide whether you want to add new sites or edit the existing ones for your vSnap servers.

Note: You can change the site name and other options for the default Primary, Secondary, and Replication sites.

Procedure

To add or edit a site, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Sites**.
2. To add new sites or edit existing sites, take the appropriate action:

Action	How to
Add a new site.	<ol style="list-style-type: none">a. Click Add site.b. Enter a site name.c. Optional: Select options to manage backup and replication operations as described in “Adding a site” on page 158. <p>Important: Modify the VMware VM allocation settings only at the direction of IBM Support.</p> <ol style="list-style-type: none">d. Click Save.
Edit a site.	<ol style="list-style-type: none">a. Select the site, and then click Edit.b. Optional: Select options to manage backup and replication operations as described in “Editing a site” on page 160. <p>Important: Modify the VMware VM allocation settings only at the direction of IBM Support.</p> <ol style="list-style-type: none">c. Click Save.

Related concepts

[“Product components” on page 5](#)

The IBM Spectrum Protect Plus solution is provided as a virtual appliance that includes storage and data movement components.

[“Managing sites” on page 158](#)

A *site* is an IBM Spectrum Protect Plus policy construct that is used to manage the placement of data in an environment.

Create backup policies

Backup policies, which are also referred to as service level agreement (SLA) policies, define parameters that are applied to backup jobs. These parameters specify the frequency of backup jobs, the retention period for backed up data, and other directives for backup operations.

Before you begin

IBM Spectrum Protect Plus includes default SLA policies as described in Chapter 8, [“Managing SLA policies for backup operations,” on page 189](#). You can use the default policies as they are or modify the policies. You can also create custom SLA policies.

For example purposes, create the following SLA policies:

An SLA policy for the resources that you want to protect

This SLA policy is applied to backup jobs for the resources that you want to protect.

An SLA policy for the IBM Spectrum Protect Plus catalog

This SLA policy is applied to backup jobs for the IBM Spectrum Protect Plus catalog. The catalog consists of data such as application configuration settings, SLA policies, backup storage settings, and information about registered resources, restore points, and jobs.

A catalog backup is helpful in a natural disaster or other unexpected event. By accessing data in the catalog backup, you can help to restore normal system operations.

To optimize backup jobs, always create separate SLA policies for backing up resources and for backing up the IBM Spectrum Protect Plus catalog. To balance the system workload, ensure that the SLA policies do not define run times that would cause the jobs to run concurrently.

About this task

For example purposes, this task describes how to create an SLA policy that defines the backup operation for VMware resources to a vSnap server and separate policy to back up the IBM Spectrum Protect Plus catalog to a vSnap server. For information about how to set up SLA policies for other configurations, see Chapter 8, “Managing SLA policies for backup operations,” on page 189.

Procedure

To create the SLA policies, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. To create SLA policies for resources and IBM Spectrum Protect Plus, take the appropriate action:

Action	How to
Create an SLA policy for resources.	<ol style="list-style-type: none"> a. Select Virtualized systems in the Category list. b. Click Tiered vSnap and then click Next. The SLA policy options are displayed on the Policy rules page. c. Complete the steps in the wizard. For assistance, see “Creating an SLA policy for hypervisor backup to a vSnap server” on page 198.
Create an SLA policy for the IBM Spectrum Protect Plus catalog.	<ol style="list-style-type: none"> a. Select IBM Spectrum Protect Plus catalog in the Category list. b. Click Catalog to vSnap and then click Next. The SLA policy options are displayed on the Policy rules page. c. Complete the steps in the wizard. For assistance, see “Creating an SLA policy for the IBM Spectrum Protect Plus catalog to back up to a vSnap server” on page 192.

Related concepts

[“Managing SLA policies for backup operations”](#) on page 189

Service level agreement (SLA) policies, also known as backup policies, define parameters for backup jobs. These parameters include the frequency and retention period of backups, the backup location, and the

option to replicate or copy backup data. You can use predefined SLA policies, or customize them to meet your needs.

Create a user account for the application administrator

Create a user account for an administrator who can run backup and restore operations for the resources in your environment.

Before you begin

For example purposes, the following steps show how to create an account for an individual user who is responsible for protecting VMware data by using vSnap server as a backup storage provider. This account uses an existing user role and resource group.

To create an account for an LDAP group, see [“Creating a user account for an LDAP group” on page 466](#).

To create custom user roles and resource groups, see [“Creating a resource group” on page 457](#) and [“Creating a role” on page 463](#)

Procedure

To create an account for an application administrator, complete the following steps:

1. In the navigation panel, click **Accounts > User**.
2. Click **Add User**. The **Add User** pane is displayed.
3. Click **Select the type of user or group you want to add > Individual new user**.
4. Enter a username and password for the application administrator.
5. Click **Add new permission**. Expand the permission pane that appears.
6. In the **Assign Role** section, click **VM Admin**.
The permissions for this role are shown in the **Permission Groups** section.
7. In the **Choose Resource Groups To Assign** section, select **All Resources**.

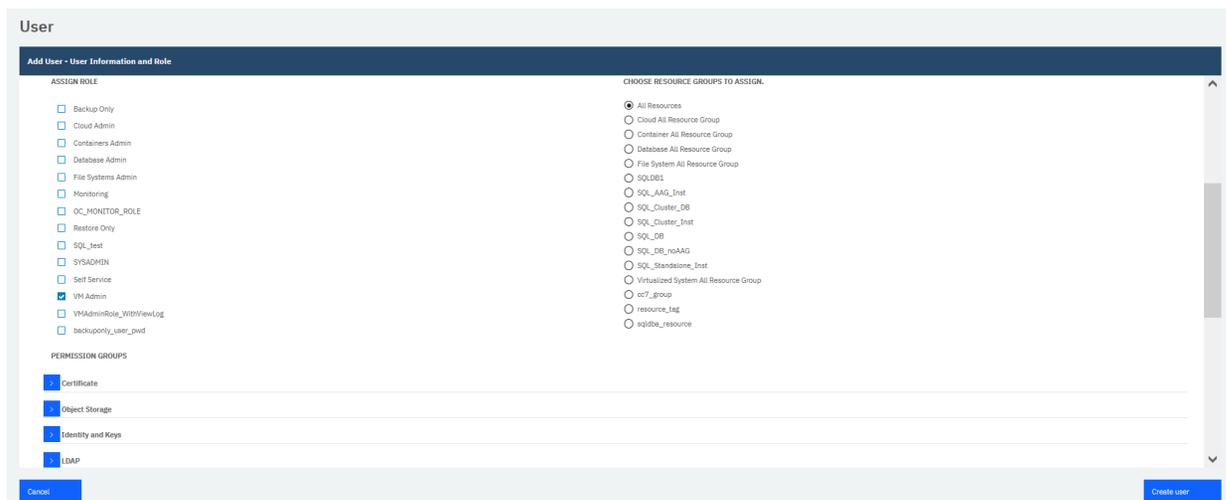


Figure 6. Creating a user account and assigning a role

8. Click **Create user**.

Related concepts

[“Managing user access” on page 455](#)

By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Back up resources

Add the resources that you want to protect to IBM Spectrum Protect Plus and create a backup job for those resources.

Add resources to protect

Resources are entities that you want to protect. After a resource is registered, an inventory of the resource is captured and added to the IBM Spectrum Protect Plus inventory, enabling you to complete backup and restore jobs, as well as to run reports.

About this task

For example purposes, this task describes how to add a VMware vCenter Server instance. To add other types of resources, see the instructions by resource type in the following sections:

- [Chapter 9, “Protecting virtualized systems,” on page 213](#)
- [Chapter 10, “Protecting file systems,” on page 267](#)
- [Chapter 11, “Protecting data on cloud systems,” on page 281](#)
- [Chapter 12, “Protecting databases,” on page 289](#)

Procedure

To add a vCenter Server instance, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. Click **Manage vCenter**, and then click **Add vCenter**.
3. Populate the fields in the **vCenter Properties** section:

Hostname/IP

Enter the resolvable IP address or a resolvable path and machine name.

Use existing user

Enable to select a previously entered user name and password for the vCenter Server instance.

Username

Enter your user name for the vCenter Server instance.

Password

Enter your password for the vCenter Server instance.

Port

Enter the communications port of the vCenter Server instance. Select the **Use TLS** checkbox to enable an encrypted Transport Layer Security (TLS) connection. The typical default port is 80 for non TLS connections or 443 for TLS connections.

4. In the **Options** section, configure the following option:

Maximum number of VMs to process concurrently per ESX server and per SLA

Set the maximum number of concurrent VM snapshots to process on the ESX server. The default setting is 3.

The following figure shows populated fields.

The screenshot shows the VMware vCenter configuration interface. At the top, there is a 'Manage vCenter' button. Below it, the 'vCenter Properties' section contains the following fields:

- Hostname/IP: 192.0.2.0
- Use existing user:
- Username: admin_192.0.2.0
- Password: masked with dots
- Port: 443
- Use TLS

The 'Options' section contains a dropdown menu for 'Maximum number of VM's to process concurrently per ESX server' with the value 3. At the bottom, there are 'Cancel' and 'Save' buttons.

Figure 7. Adding a vCenter Server instance

5. Click **Save**.

IBM Spectrum Protect Plus confirms a network connection, adds the resource to the database, and then catalogs the resource. If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to verify and possibly fix the connections.

Add resources to a job definition

Before you can back up a resource, you must create a job definition that associates the resource with one or more backup policies, also referred to as SLA policies.

About this task

For example purposes, this task describes how to select an SLA policy for resources that are in a VMware vCenter.

Procedure

To select an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. Select the resources that you want to back up. You can select all resources in a vCenter or drill down to select specific resources.

Use the search function to search for available resources and toggle the displayed resources by using the **View** filter. Available options are **VMs and Templates**, **VMs**, **Datastore**, **Tags and Categories**, and **Hosts and Clusters**. Tags, which are applied in vSphere, make it possible assign metadata to virtual machines.

Note: You must assign tags at the VM guest level for them to be utilized for backup exclusion rules based on tags or to be used as filtering for reports in IBM Spectrum Protect Plus.

The following figure shows a hard disk that is selected for backup:



Figure 8. Selecting resources for backup

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup data criteria to the job definition.

The following figure shows the SLA policy **Copper** selected:

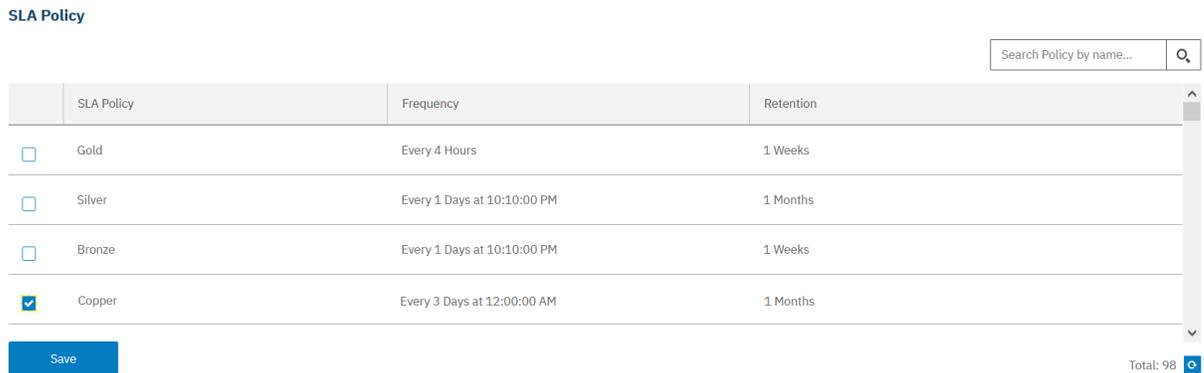


Figure 9. Selecting an SLA policy

4. To create the job definition by using default options, click **Save**.

The job name is auto generated and is constructed of the resource type followed by the SLA policy that is used for the job. For this example job, the name `vmware_Copper` is created.

5. Optional: To configure additional options, click **Select Options** and follow the instructions in [“Backing up VMware data”](#) on page 220.
6. Click **Save**.

After the job definition is saved, available virtual machine disks (VMDKs) in a virtual machine are discovered and are shown when **VMs and Templates** is selected in the **View** filter. By default, these VMDKs are assigned to the same SLA policy as the virtual machine. Optionally, to define a more granular policy by excluding individual VMDKs, follow the instructions in [“Excluding VMDKs from the SLA policy for a job”](#) on page 225.

Results

The job runs as defined by the SLA policies that you selected, or you can manually run the job by following the steps in [“Start a backup job”](#) on page 104.

Related concepts

[“Protecting IBM Spectrum Protect Plus”](#) on page 427

Protect the IBM Spectrum Protect Plus application by backing up the catalog. The catalog consists of data such as application configuration settings, SLA policies, registered resources, restore points, backup storage settings, and job information.

Back up the IBM Spectrum Protect Plus catalog

In addition to backing up the resources that you want to protect, you must also back up the IBM Spectrum Protect Plus catalog for disaster recovery scenarios.

Procedure

The catalog consists of data such as application configuration settings, SLA policies, registered resources, restore points, backup storage settings, and job information.

To back up IBM Spectrum Protect Plus catalog, complete the following steps:

1. In the navigation panel, click **Manage Protection > IBM Spectrum Protect Plus > Backup**.
2. Select the SLA policy that you created for the catalog backup job in [“Create backup policies” on page 98](#).
3. Click **Save** to create the job definition.

Results

The job runs as defined by the SLA policy that you selected, or you can manually run the job by following the steps in [“Start a backup job” on page 104](#).

Start a backup job

You can start a backup job on demand outside of the schedule that is set by the SLA policy.

Procedure

To start a backup job on demand, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**, and open the **Schedule** tab.
If your job is not a scheduled job, but is an on-demand job, click the **Job History** tab.
2. Choose the job that you want to run and click the **Start** action as shown in the following figure:

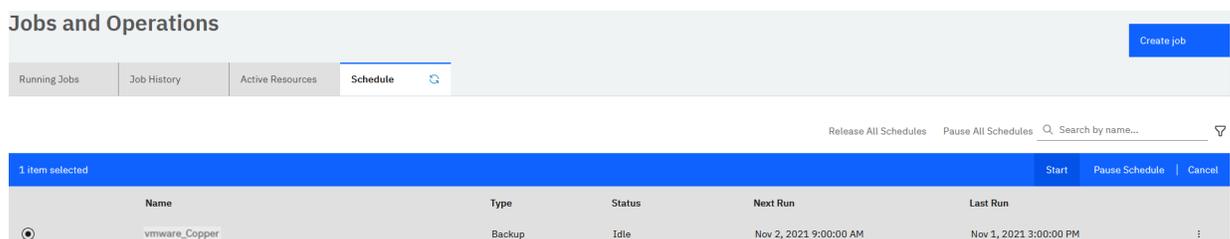


Figure 10. Starting a job

What to do next

To view the job log, click the job in the **Running Jobs** tab.

The log screen shows the following details:

- Status: shows whether the message is an error, warning, or informational message.
- Time: shows the time stamp of the message.
- ID: shows the unique identifier for the message if applicable.
- Description: shows what the message is.

You can download a job log from the page by clicking **Download .zip**. If you want to cancel the job, click **Actions > Cancel Job Type**.

Related concepts

[“Managing jobs and operations” on page 431](#)

You can manage and monitor jobs in the **Jobs and Operations** window. You can also configure scripts to run before or after jobs.

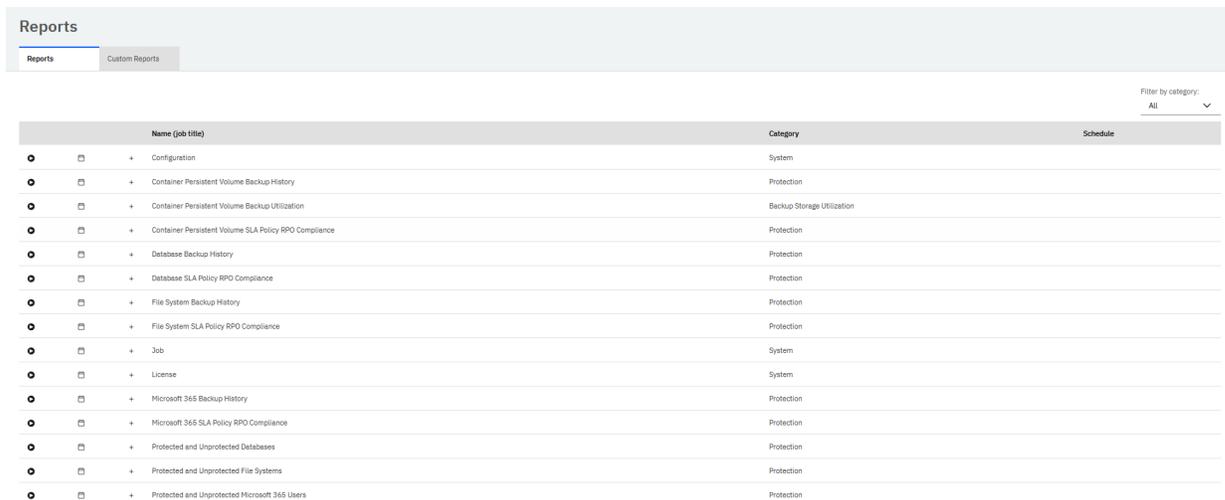
Run a report

Run reports with default parameters or custom parameters.

Procedure

To run a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click the **Reports** tab.



Name (job title)	Category	Schedule
Configuration	System	
Container Persistent Volume Backup History	Protection	
Container Persistent Volume Backup Utilization	Backup Storage Utilization	
Container Persistent Volume SLA Policy RPO Compliance	Protection	
Database Backup History	Protection	
Database SLA Policy RPO Compliance	Protection	
File System Backup History	Protection	
File System SLA Policy RPO Compliance	Protection	
Job	System	
License	System	
Microsoft 365 Backup History	Protection	
Microsoft 365 SLA Policy RPO Compliance	Protection	
Protected and Unprotected Databases	Protection	
Protected and Unprotected File Systems	Protection	
Protected and Unprotected Microsoft 365 Users	Protection	

Figure 11. Selecting a report to run

3. Run the report by clicking the **Run Report** (🎯) icon beside the report.
 - To run the report with custom parameters, set the parameters in the **Run Report** window, and click **Run**. Parameters are unique to each report.
 - To run the report with default parameters, click **Run**.

Related concepts

[“Managing reports and logs” on page 443](#)

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Chapter 7. Configuring the system environment

System management tasks include adding backup storage, managing sites, registering Lightweight Directory Access Protocol (LDAP) or Simple Mail Transfer Protocol (SMTP) servers, and managing keys and certificates for cloud resources.

Maintenance tasks include reviewing the configuration of the IBM Spectrum Protect Plus virtual appliance, collecting log files for troubleshooting, and managing Transport Layer Security (TLS) certificates.

All system management tasks apply to IBM Spectrum Protect Plus whether it is installed as a set of containers or as a virtual appliance, except for the tasks in [“Configuring IBM Spectrum Protect Plus installed as a virtual appliance”](#) on page 180. The tasks in this section apply only to IBM Spectrum Protect Plus virtual appliance installations.

Integrating with IBM Spectrum Protect

You can back up virtual machine snapshots directly to IBM Spectrum Protect directory-container storage pools by using Open Snap Store Manager (OSSM). To support direct data backup operations from IBM Spectrum Protect Plus, you must install and configure the OSSM component where the IBM Spectrum Protect server is installed.

You can monitor your IBM Spectrum Protect Plus environment from IBM Spectrum Protect Operations Center. For more information, see [“Monitoring IBM Spectrum Protect Plus from the Operations Center”](#) on page 131

Protecting VMware data using Open Snap Store Manager

To protect VMware data using Open Snap Store Manager (OSSM), you must install and configure the OSSM component on the IBM Spectrum Protect server.

You must register the Open Snap Store Manager (OSSM) storage server with IBM Spectrum Protect Plus to backup virtual machine snapshots as an alternative to using vSnap server.

Review the following restrictions and tips:

- A site defined on IBM Spectrum Protect Plus can only contain one IBM Spectrum Protect server and OSSM server. Also the vSnap storage and OSSM storage cannot be mixed in the same IBM Spectrum Protect Plus site.
- You can select one site per SLA. Selecting multiple sites for a single SLA is not supported.
- The OSSM storage server supports multiple VADP proxies in a site but can allocate only one storage per site.
- For OSSM backups, virtual machines are always placed on dedicated volumes. The global preference for grouping virtual machines does not apply to vSnap storage.
- You can replicate backup data from an OSSM storage on a primary site to as OSSM storage on a secondary site. The replication process is run by IBM Spectrum Protect server using storage rules.
- Establish replication partnership between OSSM storage hosts. You must create replication partnership from the IBM Spectrum Protect source server to IBM Spectrum Protect target server.
- The OSSM storage server does not support replication groups running concurrently. Prior to 8.1.17, there was no overlap in job schedules. Beginning with 8.1.17, the overlapping job schedules is supported. However, the secondary schedule or SLA will only start after the primary job is completed.

Installing and upgrading Open Snap Store Manager

Before you directly back up data from IBM Spectrum Protect Plus to IBM Spectrum Protect, you must install and configure the OSSM component on the IBM Spectrum Protect server. If your environment was used for the OSSM feature available in IBM Spectrum Protect 8.1.15 or previous version, you must

perform the following cleanup steps prior to upgrading OSSM to the newer version. If you are upgrading the IBM Spectrum Protect server from 8.1.16 or later, you do not need to perform the cleanup steps.

Important: If you are installing or upgrading either the IBM Spectrum Protect server or the OSSM component, ensure that both the server and the OSSM component are at the same version.

Cleaning up Open Snap Store Manager

Before you upgrade the older version of Open Snap Store Manager (OSSM) to version 8.1.17, you must clean up the previously installed OSSM. The cleanup operation removes the files that were created during the workload simulation so that you can install the newer version of OSSM.

Important:

- If you are upgrading the IBM Spectrum Protect server from 8.1.16 or later, you do not need to perform the cleanup steps.
- If you are upgrading the IBM Spectrum Protect server from 8.1.15 or earlier, you must perform the following cleanup steps prior to upgrading OSSM to the newer version.

Cleaning up OSSM to upgrade from version 8.1.15 to version 8.1.17

Before you upgrade the Open Snap Store Manager (OSSM) version 8.1.15 to version 8.1.17, clean up the previously installed OSSM.

Before you begin

Clean up IBM Spectrum Protect Plus that is used for OSSM with IBM Spectrum Protect 8.1.15 by completing the following steps:

1. Expire any restore points for the OSSM SLA.
2. Delete if any schedules created for an OSSM SLA.
3. Optionally, delete if any SLA policy created for OSSM.
4. Uninstall the VADP proxy used for the OSSM storage.
5. Delete the OSSM storage.
6. Optionally, remove the Site created for OSSM storage.
7. Run the maintenance schedule.

Procedure

To clean up IBM Spectrum Protect 8.1.15 that is used for the OSSM technology preview, complete the following steps:

Tip: No action is necessary to remove the OSSM node, file space, and domain from IBM Spectrum Protect 8.1.15 environment. This will be done automatically when you upgrade OSSM to version 8.1.17.

1. On the system that is used for your VADP Proxy, log in as a root user and stop the OSSM services by issuing the following operating system commands:

```
systemctl stop ossm
```

```
systemctl stop ossm-db
```

2. On the system that is used for your VADP Proxy, as a root user, delete the /ossm directory by issuing the following operating system command:

```
rm -r /ossm
```

3. On your IBM Spectrum Protect server, log in as a root user and stop the OSSM services by issuing the following operating system command:

```
systemctl stop ossm.service
```

4. Disable the OSSM services by issuing the following command:

```
systemctl disable ossm.service
```

5. Determine whether any containers remain undelete by issuing the following command as the IBM Spectrum Protect administrator:

```
query container stgpool=<name of storage pool used for technology preview>
```

Note: After you upgrade the version, the data extents that were created during the technology preview are typically deleted within 36 hours. However, some of the newly created containers might remain undelete.

6. If any containers are still detected in the storage pool, restart the IBM Spectrum Protect server at your convenient time to delete them.

What to do next

After cleaning up the previously installed OSSM technology preview, upgrade both the IBM Spectrum Protect server and OSSM to version 8.1.17. For instructions, see [“Upgrading Open Snap Store Manager from 8.1.15 and earlier to version 8.1.17”](#) on page 111.

Note: The IBM Spectrum Protect server and OSSM must be upgraded at the same time.

Cleaning up OSSM to upgrade from version 8.1.13 or 8.1.14 to version 8.1.17

Before you upgrade the Open Snap Store Manager (OSSM) technology preview version 8.1.13 or 8.1.14 to version 8.1.17, clean up the previously installed OSSM technology preview.

Procedure

To clean up the technology preview, complete the following steps:

1. Remove the OSSM node, file space, and domain by taking the following actions:
 - a. In the Operations Center menu, click **Technology Preview**. This opens the **Technology Preview** screen.
 - b. In the **OSSM** window, click **Reset**.

Note: You might see an exclamation point while the server and agents are reset. Typically, the reset process takes less than 1 minute. No action is required.
2. Log in as a root user and stop the OSSM services by issuing the following operating system command:

```
systemctl stop ossm.service
```

3. Disable the OSSM services by issuing the following command:

```
systemctl disable ossm.service
```

4. Uninstall the OSSM component by taking the following steps:

- a. Change to the IBM Installation Manager eclipse directory by issuing the following command:

```
cd /opt/IBM/InstallationManager/eclipse/tools
```

- b. Start IBM Installation Manager in console mode by issuing the following command:

```
./imcl -c
```

- c. Select the option to remove a previously installed software package by entering 5 and pressing Enter.
 - d. Follow the steps in the Installation wizard. Select **IBM Spectrum Protect** as the package group from where you want to uninstall a component, and then select **Open Snap Store Manager**.
5. Remove the OSSM directory, which is `<instance_user_home_directory>/ossm`. For example, if the home directory of the instance user is `/home/tsminst1`, issue the following command:

```
rm -r /home/tsminst1/.ossm
```

6. Remove the 2 GB directory that you created for the OSSM database. For example, if the directory is /home/tsminst1/ossmDatabase, issue the following command:

```
rm -r /home/tsminst1/ossmDatabase
```

7. On the IBM Spectrum Protect server, remove the administrator that you created for the technology preview. For example, if the administrator is named ADMINOSSM, you would issue the following command:

```
remove admin adminossm
```

Tip: After you reset the environment, the extents that were created for the workload simulation are typically deleted within 36 hours. However, some of the newly created containers might remain.

8. Determine whether any containers remain by issuing the following command:

```
query container stgpool=techpreview1
```

9. If any containers are still detected in TECHPREVIEW 1 storage pool, restart the IBM Spectrum Protect server at your convenient time to delete them.
10. Wait until the containers are deleted before you proceed to the next step.
11. Delete the TECHPREVIEW1 storage pool and any storage pool directories that were created. On the IBM Spectrum Protect server, complete the following steps:
 - a. Query the names of the storage pool directories that were used for the workload simulation by issuing the following command:

```
query stgpool_dir stgpool=techpreview1
```

- b. Delete each storage pool directory that was created for the workload simulation by issuing the following command:

```
delete stgpool_dir techpreview1 <stgpool_directory_name>
```

where *stgpool_directory_name* specifies the name of a storage pool directory that was used in the simulation.

- c. Delete the storage pool that was created for the workload simulation:

```
delete stgpool techpreview1
```

What to do next

After cleaning up the previously installed OSSM technology preview, upgrade the OSSM technology preview to version 8.1.17. For instructions, see [“Upgrading Open Snap Store Manager from 8.1.15 and earlier to version 8.1.17” on page 111.](#)

Upgrading Open Snap Store Manager

The Open Snap Store Manager (OSSM) can be upgraded to a newer version by using a graphical wizard or the command line in console mode.

Upgrading Open Snap Store Manager from version 8.1.16 to version 8.1.17

The Open Snap Store Manager (OSSM) version 8.1.16 can be upgraded to version 8.1.17 by using a graphical wizard or the command line in console mode.

Before you begin

You must stop the OSSM services on the IBM Spectrum Protect server by using the following command before you upgrade to version 8.1.17.

```
systemctl stop ossm
```

About this task

The instructions for upgrading OSSM are the same as the instructions for installing OSSM, with the exception that 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.

What to do next

On the IBM Spectrum Protect server, configure the OSSM instance by using a wizard. For more information, see [“Configuring the Open Snap Store Manager instance” on page 112](#).

Upgrading Open Snap Store Manager from 8.1.15 and earlier to version 8.1.17

The older versions of Open Snap Store Manager (OSSM) can be upgraded to version 8.1.17 by using a graphical wizard or the command line in console mode.

Before you begin

You must clean up the environment used for the previously installed OSSM technology preview before you upgrade to version 8.1.17. To prepare your environment for the upgrade, see [“Cleaning up Open Snap Store Manager” on page 108](#).

About this task

The instructions for upgrading OSSM are the same as the instructions for installing OSSM, with the exception that 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.

What to do next

On the IBM Spectrum Protect server, configure the OSSM instance by using a wizard. For more information, see [“Configuring the Open Snap Store Manager instance” on page 112](#).

Note: If you are upgrading an IBM Spectrum Protect server from 8.1.15 or earlier, you must run the configuration wizard. The configuration wizard may fail on the first run due to a known issue. You must re-run the wizard to complete the configuration.

Installing Open Snap Store Manager

On the IBM Spectrum Protect server, use IBM Installation Manager to install the Open Snap Store Manager (OSSM) server component. You can install OSSM by using the graphical user interface (GUI) or the command-line interface (CLI).

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 Spectrum Protect server.
- Ensure that the OSSM server and the IBM Spectrum Protect server are on the same level.
- The IBM Spectrum Protect server and the OSSM component must be installed on the same system:
 - If you are upgrading the IBM Spectrum 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 Spectrum 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

After installing OSSM, you must configure the OSSM instance. For more instructions, refer to [“Configuring the Open Snap Store Manager instance”](#) on page 112.

Configuring the Open Snap Store Manager instance

On the IBM Spectrum Protect server, use a wizard to configure the Open Snap Store Manager (OSSM) instance. The wizard guides you to configure the OSSM component by using either the graphical user interface (GUI) or the command-line interface (CLI).

Before you begin

Ensure that the following prerequisites are met:

- The OSSM component must be installed on the same system where IBM Spectrum Protect server is installed.
- Ensure that the OSSM server and the IBM Spectrum Protect server are on the same level.
- To configure the OSSM instance on the IBM Spectrum Protect server, you must have root user privileges.
- The IBM Spectrum Protect server must be running.

Procedure

To configure the OSSM instance, complete the following steps on the IBM Spectrum Protect server.

1. As a root user, on the server where IBM Spectrum Protect and the OSSM component are installed, change to the following directory:

```
<installation_directory>/ossm/bin
```

where *<installation_directory>* represents the directory where OSSM is installed.

For example, run the following command:

```
cd /opt/tivoli/tsm/ossm/bin
```

2. To start the OSSM instance configuration wizard, run the following command:

```
./ossmcfg.bin
```

This opens the **Open Snap Store Manager (OSSM) Instance Configuration Wizard** screen.

Important: By default, the OSSM instance configuration wizard runs in the GUI mode. Alternatively, to run the OSSM instance configuration wizard specifically in CLI or GUI mode, issue one of the following commands:

- To run the OSSM instance configuration wizard in console or CLI mode, issue the following command:

```
./ossmcfg.bin -i console
```

- To run the OSSM instance configuration wizard in GUI mode, issue the following command:

```
./ossmcfg.bin -i gui
```

3. Select a language of your preference from the drop-down list, and click **OK**.
4. On the **OSSM Database Configuration** screen, you can click the provided links for more information. Click **Next** to continue.
5. On the **OSSM Configuration Modes** screen, select one of the following configuration modes that you want to complete:

Important: If you are configuring the OSSM instance for the first time, you must first select the **Create a new OSSM instance** option and follow the steps to create the OSSM instance.

- [Create a new OSSM instance](#)

Important: Following are the additional options that the OSSM instance configuration wizard provides to help resolve problems that may be encountered during communication with IBM Spectrum Protect Plus, which are not part of the OSSM configuration.

- [Synchronize the IBM Spectrum Protect server administrator](#)
- [Generate a new registration key](#)
- [List API Keys](#)
- [Revoke an API Key](#)

Create a new OSSM instance

Use this option to create a new OSSM instance on the primary OSSM server.

About this task

With this option you can complete the following actions:

- Create and start the OSSM service
- Update the IBM Spectrum Protect server instance user with the OSSM server
- Create and display the newly created registration key
- If required, update the firewall rule settings

Complete the following steps to create a new OSSM instance:

Procedure

1. On the **Open Snap Store Manager Instance Configuration** screen, complete the following fields, and then click **Next**:

IBM Spectrum Protect server instance user

Select the IBM Spectrum Protect instance user ID from the drop-down list. The listed instance user IDs in the drop-down list are the same as the IDs defined with IBM Spectrum Protect server instance.

OSSM listening port

Specify the port number that is used by the OSSM server to communicate with the IBM Spectrum Protect Plus server. If the `firewalld` services are active on the system where OSSM is installed, the port is added to the firewall rules. The default port number is 3337.

Depending on the firewall service that you are using, you might need to open the ports manually. For more information, see [“Prerequisites for creating a VADP proxy” on page 119](#).

2. On the **IBM Spectrum Protect Server Connection** screen, complete the following fields, and then click **Next**:

Server address

Provide the fully qualified domain name (FQDN) of the IBM Spectrum Protect server.

Server administrative ID

Provide any server administrative ID with system level authority. The ID is not exclusive for OSSM. The control agent temporarily uses this ID to register its own administrative ID.

Server administrator password

Provide the IBM Spectrum Protect server administrator password.

Important: If the IBM Spectrum Protect server administrative ID uses the multifactor authentication (MFA), you must provide the MFA password to update the password.

Port

Specify the IBM Spectrum Protect server instance port number. The port must use the Secure Socket Layer (SSL) protocol.

3. On the **Preview** screen, review the specified settings, and then click **Configure**.

Installation begins and the **Installing OSSM Instance Configuration Wizard** screen takes few minutes to complete the installation.

4. On the **Spectrum Protect Plus Registration Key** screen, copy the OSSM registration key and save to a secure location. Click **Next** to continue.

Important: The key is required to register the OSSM instance with the IBM Spectrum Protect Plus server.

5. On the **Successful Configuration of OSSM** screen, review the OSSM service that is created and started on the IBM Spectrum Protect server, and then click **Done** to close the wizard.

What to do next

Navigate to the link provided at the end of the configuration wizard for the next steps to connect with the IBM Spectrum Protect Plus server. Before you perform the next steps, ensure that the prerequisites are met to create a VADP proxy. For instructions, see [“Prerequisites for creating a VADP proxy” on page 119](#).

Synchronize the IBM Spectrum Protect server administrator

Use this option to synchronize the IBM Spectrum Protect server administrator password with the OSSM server.

About this task

Synchronize the administrator password in the following situations:

- When you restore the IBM Spectrum Protect server database to an earlier point in time, the secret store password that is associated with the OSSM control agent does not match with the IBM Spectrum Protect server credentials.
- If the IBM-OSSM-ADMIN administrator ID is locked or its password changed, or another situation has otherwise blocked OSSM from logging in as IBM-OSSM-ADMIN.

Important: The `ossm.service` must be stopped before you perform this update.

Procedure

Complete the following steps to synchronize the IBM Spectrum Protect server administrator password with the OSSM server:

1. On the **Synchronize IBM Spectrum Protect server administrator** screen, complete the following fields, and then click **Next**:

IBM Spectrum Protect server instance user

Select the IBM Spectrum Protect instance user ID from the drop-down list. The listed instance user IDs in the drop-down list are the same as the IDs defined with IBM Spectrum Protect server instance.

Server administrative ID

Provide any server administrative ID with system level authority. The ID is not exclusive for OSSM. This ID will be used to reset the password for the IBM-OSSM-ADMIN that was created during configuration.

Server administrator password

Provide the IBM Spectrum Protect server administrator password.

Important: If the IBM Spectrum Protect server administrative ID uses the multifactor authentication (MFA), you must provide the MFA password to update the password.

2. On the **Preview** screen, review the specified settings, and then click **Configure**.

Generate a new registration key

Complete the following steps to create a new OSSM registration key for a specific IBM Spectrum Protect server instance.

Before you begin

Important: The `ossm.service` must be running.

Procedure

1. On the **Generate a New Registration Key** screen, complete the following fields, and then click **Next**:

IBM Spectrum Protect server instance user

Select the IBM Spectrum Protect instance user ID from the drop-down list. The listed instance user IDs in the drop-down list are the same as the IDs defined with IBM Spectrum Protect server instance.

IBM Spectrum Protect server instance user password

Provide the IBM Spectrum Protect server instance user password.

2. On the **Preview** screen, review the specified settings, and then click **Configure**.
3. On the **Spectrum Protect Plus Registration Key** screen, copy the OSSM registration key and save to a secure location. Click **Next** to continue.

Important: The key is required to register the OSSM instance with the IBM Spectrum Protect Plus server.

4. On the **Successfully Generated Registration Key** screen, read the information for additional actions that are required to perform on Spectrum Protect Plus server, and then click **Done** to close the wizard.

List API Keys

Use this option to list all API keys for a specific IBM Spectrum Protect server instance and to verify that the keys have valid identifiers.

Before you begin

Important: The `ossm.service` must be running.

Tip: If the OSSM instance configuration wizard is running in the console mode, and you plan to Revoke an API key, copy the appropriate key here.

Procedure

1. On the **List API Keys** screen, complete the following fields, and then click **Next**:

IBM Spectrum Protect server instance user

Select the IBM Spectrum Protect instance user ID from the drop-down list. The listed instance user IDs in the drop-down list are the same as the IDs defined with IBM Spectrum Protect server instance.

IBM Spectrum Protect server instance user password

Provide the IBM Spectrum Protect server instance user password.

2. On the **Preview** screen, review the specified settings, and then click **Configure**.
3. On the **Successfully Listed API Keys** screen, verify the API keys, and then click **Done** to close the wizard.

Revoke an API Key

If the key is not recognized, becomes inactive, lost, or compromised, use this options to revoke an API key for the specific IBM Spectrum Protect server instance.

Before you begin

Important: If the OSSM instance configuration wizard is running in the console mode, you must specify the exact key that you want to revoke. If necessary, run the wizard and choose the **List API Keys** option to discover the appropriate key, then return to the **Revoke API Key** option.

Procedure

1. On the **Revoke API Key** screen, complete the following fields, and then click **Next**:

IBM Spectrum Protect server instance user

Select the IBM Spectrum Protect instance user ID from the drop-down list. The listed instance user IDs in the drop-down list are the same as the IDs defined with IBM Spectrum Protect server instance.

IBM Spectrum Protect server instance user password

Provide the IBM Spectrum Protect server instance user password, and then click **Generate List of Existing Keys**.

Important: If the registration keys are not generated, ensure that the entered password is correct. For more information on the output from the commands, find the details in the `/var/tivoli/tsm/ossmcfg.trc` file.

Registration Keys

Select the key to revoke from the drop-down list.

2. On the **Preview** screen, review the specified settings, and then click **Configure**.
3. On the **Successfully Revoked API Key** screen, verify the revoked key, and then click **Done** to close the wizard.

Registering OSSM administrator on the IBM Spectrum Protect server

On the IBM Spectrum Protect server, control agent registers and manages its own administrator ID for OSSM that is named as IBM-OSSM-ADMIN.

To register the OSSM administrator ID for the control agent, command line option is issued under the following circumstances:

- When you set up the control agent manually for the first time.
- While performing the following operations:
 - Database restore to an earlier point in time.
 - If the administrator notices errors on the IBM Spectrum Protect Plus server while attempting to perform backups or administrative actions with OSSM storage.

Investigate the IBM Spectrum Protect activity log for messages such as administrator locked and invalid sign on attempt. The messages include the administrator ID IBM-OSSM-ADMIN. To resolve the problems, you must issue the command line option and then the control agent registers its own administrator ID.

Important:

- If you use the configuration wizard, the wizard runs the command automatically for you in problematic situations and resolve problems.
- Though you can set up control agent manually, to use the OSSM feature more efficiently you can use the configuration wizard to set the control agent for the first time. The configuration wizard runs the command automatically for you to register the OSSM administrator ID.
- The control agent must not be running when the command is issued. To check the current running state of control agent, issue the following command:

```
systemctl status ossm
```

After issuing this command, if you see the status is active and running, issue the following command to stop the control agent:

```
systemctl stop ossm
```

The following is an example of the command that you might need to issue or the configuration wizard runs:

Important: The `ossm systemd` service must not be running when the control agent synchronizes the IBM-OSSM-ADMIN administrator.

```
ossmctl service sync-admin -a sp_server_admin -c /path/to/ossm/config_primary.json
```

Where:

- `sp_server_admin` is the name of the administrator's ID that you want to use to synchronize the OSSM administrator.
- `/path/to/ossm/config_primary.json` is the directory to the OSSM configuration file. The control agent uses the file to locate the information that is required to connect to the IBM Spectrum Protect server and find its secret store.

After running this command, the control agent prompts for the administrator's password. The password is required for the control agent to sign in to the server and complete the registration.

Credential requirements for the OSSM administrator

The administrator credential that is used by the control agent to define the OSSM administrator must meet the following requirements.

- If you are configuring OSSM for the first time or the IBM-OSSM-ADMIN administrator already exists and you are running a repair action in unexpected events, following are the requirements for the IBM Spectrum Protect administrator:
 - The administrator must have system level authority.
 - New sessions for administrators must not be disabled.
 - The administrator's ID must not be locked.
 - The administrator's password must not be expired.
 - The administrator authentication may be local or LDAP.
 - The administrator may have MFA enabled.
- If the password is MFA enabled, the one-time password must be specified when the control agent prompts for the password. The password will not be echoed as the user enters the password. Optionally, if you are using the command manually, you can echo the password by issuing the command that is similar to the following example:

```
echo "passw0rd" | service sync-admin -a myadmin -c /home/tsminst1/ossm/config_primary.json
```

The administrator will see the prompt for the password but password is accepted from the echo command instead of waiting for user's input.

- The **ossm systemd** service must not be running when the control agent runs a command.

Establishing communication between two IBM Spectrum Protect servers

If you want to enable replication for backup data from IBM Spectrum Protect Plus, you must establish server-to-server communication between two IBM Spectrum Protect servers.

Procedure

To establish communication between two IBM Spectrum Protect servers, complete the following steps by using the command builder in Operations Center or as a system administrator on both the servers:

1. On both the source and target servers, issue the following commands and set the values:

```
SET SERVERNAME
```

```
SET SERVERPASSWORD
```

```
SET SERVERHLA
```

2. You can verify the values by issuing the following command:

```
QUERY STATUS
```

3. On the IBM Spectrum Protect source server, define the IBM Spectrum Protect target server and ping server to verify communication by issuing the following command:

```
DEFINE SERVER <target_SERVERNAME> SERVERPAssword=<target_SERVERPASSWORD>  
HLA=<target_SERVERHLA> LLA=<target_TCPport>
```

```
PING SERVER <target_SERVERNAME>
```

4. On the IBM Spectrum Protect target server, define the IBM Spectrum Protect source server and ping server to verify communication by issuing the following command:

```
DEFINE SERVER <source_SERVERNAME> SERVERPAssword=<source_SERVERPASSWORD>  
HLA=<source_SERVERHLA> LLA=<source_TCPport>
```

```
PING SERVER <source_SERVERNAME>
```

Tip:

- The HLA parameter represents a hostname or IP address, as specified in the *SET SERVERHLA* command.
- The LLA parameter represents the low-level port address, as specified in the *TCPport* option in the *dmserv.opt* file and displayed in the *QUERY STATUS* output.

Prerequisites for creating a VADP proxy

In IBM Spectrum Protect Plus, running VM backup jobs through VADP requires significant system resources. By creating VADP backup job proxies, you enable load sharing and load balancing for IBM Spectrum Protect Plus backup jobs. If proxies exist, the entire processing load is shifted from the IBM Spectrum Protect Plus server onto the proxies. All system requirements and prerequisites must be met before you start creating a VADP proxy for Open Snap Store Manager (OSSM).

For the information about system requirements of a OSSM server and the VADP proxy, see [technote 6837823](#).

The following are the prerequisites for creating the VADP proxies when using Open Snap Store Manager (OSSM) as a storage server on a Linux system:

- Install the fuse package for your environment by issuing one of the following commands:
 - For Red Hat Enterprise Linux (RHEL) 8

```
dnf install fuse
```

- For SUSE Linux Enterprise Server (SLES) 15

```
zypper install fuse
```

- Make sure that `user_allow_other` is uncommented in the `/etc/fuse.conf` file.
- You must have passwordless sudo privileges to create a VADP proxy on an OSSM storage server. To create a new user or a new user group with passwordless sudo privileges, complete the following steps:
 1. Create a new user or a new user group by issuing the following command:

```
groupadd <NEWGROUP>
useradd -g <NEWGROUP> -m -d /home/<NEWUSER> -s /bin/bash <NEWUSER>
```

where `<NEWUSER>` is the username, and `<NEWGROUP>` is the user group.

2. Set a password for the new user by issuing the following command:

```
passwd <NEWUSER>
```

3. Open sudoers file by issuing the following command:

```
$ sudo visudo
```

4. Add the following line at the end of the sudoers file and save.

```
<NEWUSER> ALL=(ALL) NOPASSWD: ALL
```

where `<NEWUSER>` is the username.

- Ensure that SELinux is disabled. To disable SELinux, complete the following steps:
 1. Open the `/etc/selinux/config` file.
 2. Locate the line: `SELINUX=enforcing`.
 3. Change the value to `SELINUX=disabled`.
 4. Reboot the operating system.
- Ensure that the date-time is in sync between the VADP proxy system and the system that is hosting the IBM Spectrum Protect server, and the OSSM primary control agent service.

Enable the Network Time Protocol (NTP) service to synchronize the system time. To start automatic time synchronization with a remote NTP server, issue the following command:

```
timedatectl set-ntp true
```

Tip: Ensure that the time zone is same for both the VADP proxy and the IBM Spectrum Protect server.

- The following ports must be opened on the firewall:
 - Port 3337: OSSM control agent port
 - Port 111, 2049, and 20048: NFS mount access for certain VM restore scenarios

Depending on the firewall service that you are using, you might need to open the ports by completing one of the following steps:

- If you have the firewalld service running, all required ports are opened except 3337. You must close port 3338 and manually open port 3337.

1. Issue the following command to close the port 3338:

```
firewall-cmd --zone=public --remove-port=3338/tcp --permanent
```

2. Issue the following command to manually open the port 3337:

```
firewall-cmd --zone=public --add-port=3337/tcp --permanent
```

3. Issue the following command to activate the changes:

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

- If you have another service running than firewalld, you must manually open all the ports.

Issue the following command to open the ports:

```
firewall-cmd --zone=public --add-port=3337/tcp --permanent
firewall-cmd --zone=public --permanent --add-port=111/tcp
firewall-cmd --zone=public --permanent --add-port=2049/tcp
firewall-cmd --zone=public --permanent --add-port=20048/tcp
firewall-cmd --reload
```

- If you do not have the firewall service running, opening of ports is not required.

What to do next

To add OSSM storage server for backing up VMware data, see [“Adding the Open Snap Store Manager server as a backup storage provider”](#) on page 120.

Backing up VMware data using Open Snap Store Manager

You can run a backup job to copy the snapshots directly to the IBM Spectrum Protect by using Open Snap Store Manager (OSSM).

You must register the Open Snap Store Manager (OSSM) storage server with IBM Spectrum Protect Plus to backup virtual machine snapshots directly to IBM Spectrum Protect for long-term retention and archiving without requiring an intervening vSnap server.

Adding the Open Snap Store Manager server as a backup storage provider

You can add the Open Snap Store Manager (OSSM) as a backup storage for VMware workloads.

Before you begin

- Ensure that you have configured the OSSM instance on the IBM Spectrum Protect server. For instructions, see [“Configuring the Open Snap Store Manager instance”](#) on page 112.
- Ensure that you have the registration key that you copied during the OSSM configuration.
- Ensure that you have an empty directory-container storage pool configured on the IBM Spectrum Protect server.

- You must have a unique site that is not associated with any other OSSM storage server or any vSnap servers. For instructions about how to add a site, see [“Adding a site” on page 158](#).

Procedure

To add the OSSM as a backup storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage**.
2. Click **OSSM** to open the OSSM wizard.
3. Click **Add OSSM**. The Add OSSM Storage wizard opens.
4. Complete the following fields on the **OSSM details** page, and then click **Next**.

IBM Spectrum Protect server Hostname/IP

Enter the Hostname or IP address of the IBM Spectrum Protect server where the configured OSSM primary control agent service is also running.

Site

Select a site that you created for the OSSM storage server.

Port

Specify the port number for the control agent.

Registration Key

Enter the registration key that you copied from the IBM Spectrum Protect server OSSM configuration wizard.

Certificate

Locate the OSSM certificate `certificate.pem` in the OSSM instance configuration directory on the IBM Spectrum Protect server. By default, the OSSM certificate is stored in the following directory:

```
/home/tsminst1/ossm/.secret/
```

where `tsminst1` is the IBM Spectrum Protect instance user ID.

Note: If you are unsure of the OSSM configuration directory location, contact your IBM Spectrum Protect administrator.

In the **Certificate** field, select one of the following options to import the certificate:

Upload

- a. Copy the `/home/tsminst1/ossm/.secret/certificate.pem` file from the IBM Spectrum Protect server to the local machine where you are running the IBM Spectrum Protect Plus browser.
- b. Rename the certificate file. The name must be unique and must not match any of the existing file names.
- c. Click **Choose file** and search for the downloaded certificate in your system.
- d. Click **Upload**. This option is the default.

Copy and paste

Enter a unique name for the certificate, such as `hostname_certificate.pem`. Then, paste the contents of the certificate in the **Copy and paste certificate here** field and click **Create**.

Note: Ensure that the certificate name is unique and does not match any of the existing file names.

Use existing certificate

Click **Choose file** to select an existing certificate from the **Select a certificate** list.

5. On the **Select pool** wizard, select a specific directory-container storage pool that is used as a target storage pool, and then click **Next**.
6. On the **Review** page, review your selections, and then click **Submit**.

Results

The OSSM storage server is registered.

What to do next

After you register the OSSM storage server, create a VADP proxy to run the VMware backup. For instructions, see [“Creating VADP proxies for OSSM” on page 123](#).

Related tasks

[“Editing the Open Snap Store Manager” on page 122](#)

You can edit the options for Open Snap Store Manager (OSSM) that you registered as a backup storage for VMware workloads.

Editing the Open Snap Store Manager

You can edit the options for Open Snap Store Manager (OSSM) that you registered as a backup storage for VMware workloads.

Procedure

To edit the options in the OSSM backup storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage**.
2. Click **OSSM** to open the OSSM wizard.
3. Click the OSSM storage server that you registered, and then click **Edit**.
4. Edit the following fields:

Port

Specify the port number for the control agent.

Registration Key

Enter the registration key that you copied from the OSSM configuration wizard.

Restriction: You cannot change the **IBM Spectrum Protect server Hostname/IP** and **Site** details.

Certificate

Locate the OSSM certificate `certificate.pem` in the OSSM instance configuration directory on the IBM Spectrum Protect server. By default, the OSSM certificate is stored in the following directory:

```
/home/tsminst1/ossm/.secret/
```

In the **Certificate** field, select one of the following options to import the certificate:

Upload

- a. Copy the `/home/tsminst1/ossm/.secret/certificate.pem` file from the IBM Spectrum Protect server to the local machine where you are running the IBM Spectrum Protect Plus browser.
- b. Rename the certificate file. The name must be unique and must not match any of the existing file names.
- c. Click **Choose file** and search for the downloaded certificate in your system.
- d. Click **Upload**. This option is the default.

Copy and paste

Enter a name for the certificate, such as `certificate.pem`. Then, paste the contents of the certificate in the **Copy and paste certificate here** field and click **Create**.

Use existing certificate

Click **Choose file** to select an existing certificate from the **Select a certificate** list.

5. Click **Save** to update the OSSM storage server settings.

Managing VADP backup proxies

In IBM Spectrum Protect Plus, you must create proxies to run VMware backup jobs by using vStorage API for Data Protection (VADP) in Linux environments. The proxies reduce demand on system resources by enabling load sharing and load balancing.

At least one VADP proxy must be enabled in the backup site that is specified in the SLA for VMware backups. For more information, see [“Creating VADP proxies for OSSM” on page 123](#).

The processing load is shifted off the host system and onto the proxies for VMware backup jobs. When more than one VADP proxy exists, throttling ensures that multiple proxies are optimally utilized to maximize data throughput. For each VMware virtual machine being backed up, IBM Spectrum Protect Plus determines which VADP proxy is the least busy and has the most available memory and free tasks.

If a proxy server goes down or is otherwise unavailable before the start of the job, the other proxies take over and the job is complete. If a proxy server becomes unavailable when a job is running, the job may fail.

Transport modes describe the method by which a VADP proxy moves data. The transport mode is set as a property of the proxy. Most backup and recovery jobs are later configured to use one or more proxies.

Creating VADP proxies for OSSM

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Before you begin

Restriction: For running the steps to create VADP proxies, ensure that you have a user ID with the SYSADMIN role assigned. For more information about roles, see [Managing roles](#).

Ensure that the prerequisites for creating the VADP proxy are met. For more information, see [“Prerequisites for creating a VADP proxy” on page 119](#).

For the information about system requirements of a OSSM server and the VADP proxy, see [technote 6837823](#).

Procedure

To create VMware VADP proxies, complete the following steps:

1. In the navigation panel, click **System Configuration > VADP Proxy**.
2. Click **Register Proxy**.
3. Complete the following fields in the **Install VADP Proxy** pane:

Hostname/IP

Enter the resolvable IP address or a resolvable path and machine name.

Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

Select a site

Select the unique site that is created for OSSM to associate with the proxy.

Use existing user

Enable to select a previously entered user name and password for the provider.

Important: Do not select this check box if this is your first time registering the VADP proxy for OSSM. Enter the Username and Password as described below.

Username

Enter the user name that has the passwordless sudo privileges, which you created for VADP proxy.

Password

Enter the password that you defined for VADP proxy.

4. Click **Next**.

5. Click **Submit** on the **Review** screen.

Notice: Wait for a few minutes to complete the VADP proxy installation.

6. Change the transport mode by completing the following steps:

a) On the VADP proxy screen, click the VADP proxy, which then displays the information in the adjacent details panel.

b) In the **Proxy Details** panel, click the ellipses icon (***) to open the actions menu, and then choose **Proxy Options**. The **Set VADP Proxy Options** window opens.

c) From the drop-down list of **Transport Modes**, select the **nbdssl:nbd** check box.

d) Click **Save** to close the window.

7. Repeat the previous steps for each proxy that you want to create.

Results

The proxy is added to the **VADP Proxy** table.

The VADP proxy table displays that the proxy is OSSM Enabled. To customize the columns that are displayed on each VADP proxy, click the settings icon , and then select the required columns.

What to do next

After you create a VADP proxy, associate the OSSM storage server with an SLA policy that is used for the backup job. For instructions, see [“Creating an SLA policy for VMware backup to the OSSM storage server” on page 127](#)

Cleaning up a VADP proxy created for Open Snap Store Manager

If you no longer need to use a VADP proxy that is created for Open Snap Store Manager (OSSM), you can remove the code for VADP proxy and OSSM from the IBM Spectrum Protect Plus environment.

Procedure

To cleanup a VADP proxy from your IBM Spectrum Protect Plus, complete the following steps:

1. As a root user, log in to the VADP proxy server and stop the OSSM services by issuing the following operating system commands:

```
systemctl stop ossm
systemctl stop ossm-db
```

2. Disable the OSSM services by issuing the following commands:

```
systemctl disable ossm
systemctl disable ossm-db
```

3. Uninstall the OSSM service by issuing the following command:

```
/ossm/ossctl service uninstall
```

4. Remove the OSSM directory /ossm by issuing the following command:

```
rm -r /ossm
```

5. Log in to the IBM Spectrum Protect Plus server with an admin ID that has sysadmin authority.

6. In the navigation panel, click on **System Configuration > VADP Proxy**.

7. Click the VADP proxy that you want to uninstall, which then displays the information in the adjacent details panel.

8. Click the ellipses icon **⋮** in the details panel, and then select **Uninstall**.

Disassociating a VADP Proxy

You can disassociate a VADP proxy by using the **curl** command from the IBM Spectrum Protect server.

Before you begin

- Ensure that you have the login credentials for the IBM Spectrum Protect server instance user on the system where the IBM Spectrum Protect server is installed.
- All commands will be issued on the IBM Spectrum Protect server as the IBM Spectrum Protect server instance user.
- Ensure that the **curl** command line tool is installed on the IBM Spectrum Protect server.
- The IBM Spectrum Protect server and Open Snap Store Manager (OSSM) must be running.

Procedure

1. To get the API key and the secret parameter value, open a command line and issue the following command:

```
curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -u "username:password" -XGET https://sphost:port/api/v1beta/internal/registration/key
```

where:

username specifies an instant user's ID.

password specifies an instant user's password.

sphost specifies the hostname or IP address of the system where the IBM Spectrum Protect server is running.

port specifies the port number of the OSSM central agent. The default port number is 3337.

For example:

```
curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -u "tsminst1:tsminst1" -XGET https://hostname.mycompany.com:3337/api/v1beta/internal/registration/key
```

The output is similar to the following response:

```
{
  "registration_keys": {
  },
  "api_keys": {
    "zLdHPwhYIWAPXSIz": {
      "identifier": "12345678",

```

```

    "key": "zLdHPwhYIWAPXSIz",
    "secret": "DM5RQVRgE04q0Fbsm6nHTHGXeIMx/jtjt1tC6Gnm",
    "created": "2022-05-24T21:25:24.364023582Z",
    "revoked": "0001-01-01T00:00:00Z",
    "hostname": "fake.host.com"
  }
}
}

```

Save the **key** string value and the **secret** string value. These values will be required to perform the next step.

2. To get a token for the REST requests, issue the following command:

```

curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -u
"key:secret" -XPOST https://sphost:port/api/v1beta/auth

```

where:

key specifies the key string value that you obtained from [step 1](#).

secret specifies the secret string value that you obtained from [step 1](#).

sphost specifies the hostname or IP address of the system where the IBM Spectrum Protect server is running.

port specifies the port number of the OSSM central agent. The default port number is 3337.

For example:

```

curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/
json" -u "zLdHPwhYIWAPXSIz:DM5RQVRgE04q0Fbsm6nHTHGXeIMx/jtjt1tC6Gnm" -XPOST https://
hostname.mycompany.com:3337/api/v1beta/auth

```

The output is similar to the following response:

```

{
  "Claims": {
    "aud": "12345678",
    "exp": 1654293288,
    "iat": 1654206888,
    "iss": "OSSM Control Agent",
    "nbf": 1654206888,
    "sub": "zLdHPwhYIWAPXSIz"
  },
  "token_string":
  "eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJhdWQiOiIiXmM0NTY3OCIsImV4cCI6MTY1NDI5MzI4OiwiaWF0IjoxNjU0MjA2ODg4LCJpc3MiOiJPU1NNIENvbnRyb2wgQWdlbnQiLCJyYmYiOiJlE2NTQyMDY4ODgsInN1YiI6InpMZEhQd2hZSVdBUFhTSXoifQ.sfkHJ0iCCbyY4Wd0gvi6wwqk1NW6oIeejvMUkyOm5rc"
}

```

Save the **token_string** value. The value will be required to perform the next step.

3. To get the ID number of the VADP proxy agent, issue the following command:

```

curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -H
"Authorization: Bearer token" -XGET https://sphost:port/api/v1beta/proxy

```

where:

token specifies the token_string value that you obtained from [step 2](#).

sphost specifies the hostname or IP address of the system IBM Spectrum Protect server.

port specifies the port number of the OSSM central agent. The default port number is 3337.

For example:

```

curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -H
"Authorization: Bearer
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJhdWQiOiIiXmM0NTY3OCIsImV4cCI6MTY1NDI5MzI4OiwiaWF0IjoxNjU0MjA2ODg4LCJpc3MiOiJPU1NNIENvbnRyb2wgQWdlbnQiLCJyYmYiOiJlE2NTQyMDY4ODgsInN1YiI6InpMZEhQd2hZSVdBUFhTSXoifQ.sfkHJ0iCCbyY4Wd0gvi6wwqk1NW6oIeejvMUkyOm5rc"

```

```
dBUFHtSXoifQ.sfkHJ0iCCbyY4Wd0gvi6wvqk1NW6oIeejvMUky0m5rc" -XGET https://
hostname.mycompany.com:3337/api/v1beta/proxy
```

The output is similar to the following response:

```
[
  {
    "id": 1,
    "name": "vcloud833",
    "addr": "1.23.45.678",
    "port": 3338,
    "status": 0,
    "created_time": "2022-06-02T10:13:45.285916085-07:00",
    "last_sync_time": "1999-11-30T00:00:00-07:00",
    "log_path": "",
    "version": "",
    "node": null,
    "offline_ma": false,
    "out_of_sync": false
  }
]
```

Save the proxy ID that you want to disassociate.

4. To disassociate a VADP proxy, issue the following command:

```
curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -H
"Authorization: Bearer token" -XDELETE https:// sphost:port/api/v1beta/proxy/id
```

where:

token specifies the `token_string` value that you obtained from [step 2](#).

sphost specifies the hostname or IP address of the system where the IBM Spectrum Protect server is running.

port specifies the port number of the OSSM central agent. The default port number is 3337.

id specifies the proxy ID that you obtained from [step 3](#).

For example:

```
curl -v --insecure -H 'Accept: application/json' -H "Content-Type: application/json" -H
"Authorization: Bearer
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJhdWQiOiIiXmMjM0NTY3OClzImV4cCI6MTY1NDI5MzI4OClzIjox
NjU0MjA2ODg4LCJpc3MiOiJPU1NNIENvbnRyb2wgQWdlbnQiLCJyYmYiOiJlMjNTQyMDY4ODgsInN1YiI6InpMZEhQd2hZS
V
dBUFHtSXoifQ.sfkHJ0iCCbyY4Wd0gvi6wvqk1NW6oIeejvMUky0m5rc" -XDELETE https://
hostname.mycompany.com:3337/api/v1beta/proxy/1
```

5. Repeat [step 3](#) to verify that the VADP proxy that you intended to disassociate is no longer visible in the output.

Creating an SLA policy for VMware backup to the OSSM storage server

You can create the custom SLA policies that enable you to back up VMware data to the Open Snap Store Manager (OSSM) storage server.

Before you begin

Ensure that the storage system that you want to associate with an SLA policy is configured in **System Configuration > Storage > OSSM**. For information about adding the OSSM storage server, see [“Adding the Open Snap Store Manager server as a backup storage provider”](#) on page 120.

About this task

If a virtual machine is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To create an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **Virtualized systems** in the **Category** list.
4. Click **OSSM** and then click **Next**.

The SLA policy options are displayed on the **Policy rules** page.

5. Complete the following options on the page, and then click **Next**.
 - a) In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
 - b) Optional: Select **Disable all Schedules** checkbox to disable all scheduling options in the policy.
 - c) In the **Backup Policy** section, set the following options for backup operations.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for backup operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you might select one or more days of the week. The **Start Time** applies to the selected days of the week.

Start Time

Enter the date and time when you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for backing up the data.

Only sites that are associated with an OSSM storage server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with an OSSM storage server, are not shown.

- d) In the **Replication Policy** section, set the following options to enable asynchronous replication from one OSSM server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in [“Configuring backup storage partners for Open Snap Store Manager” on page 129](#).

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this checkbox to create the replication relationship without defining a frequency or start time.

Repeats

Enter a frequency for replication operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for replicating data.

A site can only contain one OSSM storage server.

Only sites that are associated with a OSSM server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a OSSM server, are not shown.

Same retention as source selection

Select this option to use the same retention policy as the source OSSM server. To set a different retention policy, clear this option and set a different policy.

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, refer to [“Managing Jobs for VMware backups” on page 129](#) to run a backup and restore jobs.

Related information

[Editing an SLA policy](#)

[Deleting an SLA policy](#)

Managing Jobs for VMware backups

You can run a backup job and a restore job to copy the snapshots directly to the IBM Spectrum Protect server by using Open Snap Store Manager (OSSM).

- To run a backup job to the OSSM storage server for VMware, see [“Backing up VMware data” on page 220](#).
- To run a restore job from the OSSM storage server for VMware, see [“Restoring VMware data” on page 232](#).
- To monitor the backup and restore job status from the SPP dashboard, see [Chapter 14, “Managing jobs and operations,” on page 431](#).

Configuring backup storage partners for Open Snap Store Manager

You can configure your backup storage primary and secondary sites to establish replication partnerships with other sites to extend your environment. After you configure replication partners, you can copy data from one site to another for an added layer of data protection.

About this task

To add partners to a server in your storage environment, complete the following steps:

Procedure

1. In the navigation panel, click **System Configuration > Storage > OSSM**.
2. Select the OSSM server that you want to configure, and then click **Manage**.
3. Open the **Storage Partners** tab.
4. Click **Add** to add a partner to your primary or secondary backup storage host.

Storage efficiency	Storage Partners	Active Directory	Storage disk	Network interface controllers	Streaming, retrieval, and allocation rules
Configure Storage Partners					
Current partners					
Hostname/IP	Created				
knob2.tucson.ibm.com	Jul 20, 2020 12:03:36 PM				
					Remove
Available partners					
You can add the hosts below to current partners. Only compatible hosts are shown					
Hostname/IP and Pool	Capacity	Efficiency Settings			
knob9.tucson.ibm.com Site: Replication	Capacity: available of	<input checked="" type="radio"/> Compression	<input checked="" type="radio"/> Deduplication	<input checked="" type="radio"/> Encryption	Add
knob7.tucson.ibm.com Site: Primary	Capacity: available of	<input type="radio"/> Compression	<input type="radio"/> Deduplication	<input type="radio"/> Encryption	Add
usc-18.tucson.ibm.com Site: Primary	Capacity: available of	<input checked="" type="radio"/> Compression	<input type="radio"/> Deduplication		Add

What to do next

After you establish the replication partnership, see [“Creating an SLA policy for VMware backup to the OSSM storage server”](#) on page 127.

Replicating backup storage data by using Open Snap Store Manager

When you enable cross-replication in the IBM Spectrum Protect Plus environment, backup data is asynchronously replicated between two IBM Spectrum Protect servers by using Open Snap Store Manager (OSSM). For example, you can replicate backup data from the OSSM storage on a primary site to the OSSM storage on a secondary site and in reverse direction. The data replication is done by IBM Spectrum Protect server by using replication storage rules.

Before you begin

Before you replicate backup data from IBM Spectrum Protect Plus to IBM Spectrum Protect, ensure to complete the following actions on the IBM Spectrum Protect environment:

1. Ensure that you have installed and configured two IBM Spectrum Protect servers to cross-replicate the data.
2. Define server-to-server communication on both the servers. For more information, see [“Establishing communication between two IBM Spectrum Protect servers”](#) on page 118.
3. Install OSSM and configure the OSSM services on both the servers.

Restriction: By using OSSM, replication feature is supported for only cross-replication between two IBM Spectrum Protect servers and not supported for multi-target replication.

Procedure

1. Add sites in IBM Spectrum Protect Plus for both the IBM Spectrum Protect servers. You must have a unique site that is not associated with any other OSSM storage server or any vSnap servers. Each site requires use of at least one VADP proxy. For instructions about how to add a site, see [“Adding a site”](#) on page 158.
2. Add OSSM storage for both the servers. For instructions, see [“Adding the Open Snap Store Manager server as a backup storage provider”](#) on page 120.

Important: If the OSSM storage has been upgraded from IBM Spectrum Protect Plus 10.1.11 and does not have a certificate, delete the device and re-add the OSSM storage server. For instructions, see [“Adding the Open Snap Store Manager server as a backup storage provider”](#) on page 120.

3. Configure storage partner for each OSSM storage server. For the source server, add the target server as a storage partner and vice versa. For instructions, see [“Configuring backup storage partners for Open Snap Store Manager”](#) on page 129.

4. Create SLA policy for both the servers. For instructions, see “[Creating an SLA policy for VMware backup to the OSSM storage server](#)” on page 127. To select the target site for both the server SLA, complete the following actions in the navigation panel:
 - For source server SLA policy, click **Replication Policy > Backup Storage Replication** in the navigation menu and select target server for **Target Site**.
 - For target server SLA policy, click **Replication Policy > Backup Storage Replication** in the navigation menu and select source server for **Target Site**.
5. Select and assign VMs to both the server SLAs. You can also assign the same VM to both the server SLAs.

Important: If you assign the same VM to both the server SLAs, the servers cannot be scheduled to back up at the same time.
6. Run the schedule to back up VMs to both the server SLAs.
7. To run the replication manually after the first scheduled replication, in the navigation panel, go to **Jobs & Operations > Schedule** > choose the replication schedule, and then click the icon  > **Start > SLA Policy** > select **Replicate** from the drop-down list.
8. To run the cross-replication, start the **SLA Policy** for replication on both the servers at the same time.

Monitoring IBM Spectrum Protect Plus from the Operations Center

You can monitor your IBM Spectrum Protect Plus environment from IBM Spectrum Protect Operations Center. For convenience, you can also access the Operations Center directly from IBM Spectrum Protect Plus.

The Operations Center includes a dashboard for IBM Spectrum Protect Plus that provides the following information:

- A summary of job activities for a selected time period. You can view the percentages of backup, restore, and other jobs that succeeded and failed. From this summary information, you can go to more detailed information for each job type.
- A summary of the capacity and availability of vSnap servers. You can view the total disk capacity that is available to the IBM Spectrum Protect Plus server through all vSnap servers. You can also view the available capacity for each vSnap server.
- A summary of service level agreement (SLA) policies that are defined on the IBM Spectrum Protect Plus server. You can view the number of policies that have associated backup jobs. You can also view the percentage of resources that are protected by backup jobs, and the number of resources that are not protected. From this summary information, you can go to more detailed policy information.

To enable this feature, a system administrator must add the IBM Spectrum Protect Plus server to the Operations Center.

Access the Operations Center from the IBM Spectrum Protect Plus GUI

To access the Operations Center from IBM Spectrum Protect Plus, a system administrator must add the Operations Center URL on the **Global Preferences** page of the IBM Spectrum Protect Plus GUI.

You can then access the Operations Center from the IBM Spectrum Protect icon  on the menu bar.

Adding IBM Spectrum Protect Plus to the Operations Center

When you add an IBM Spectrum Protect Plus server to the Operations Center, you establish a connection between the server and the Operations Center. After this connection is established, you can use the Operations Center to monitor the IBM Spectrum Protect Plus environment.

Before you begin

Ensure that you have the URL for the Operations Center and user credentials to log on.

Procedure

To add an IBM Spectrum Protect Plus server to the Operations Center, complete the following steps:

1. On the Operations Center menu bar, click **Overviews > Protect Plus** and take one of the following actions to open the **Add Server** wizard:

Current configuration	Action
No IBM Spectrum Protect Plus servers are connected to the Operations Center.	A message indicates that no IBM Spectrum Protect Plus servers are configured. Click +Add Server .
One or more IBM Spectrum Protect Plus servers are connected to the Operations Center.	The IBM Spectrum Protect Plus dashboard is displayed. From the list of servers on the monitoring dashboard, select +Add Server . 

2. To add the IBM Spectrum Protect Plus server, follow the directions in the wizard.

On the **Authorization** page of the wizard, you are prompted to specify user credentials to access and monitor the IBM Spectrum Protect Plus server. If you have an IBM Spectrum Protect Plus account whose credentials match the Operations Center credentials, you can use that account. If you don't have matching credentials, you must create an account.

Use Operations Center credentials

Select this option to use an existing IBM Spectrum Protect Plus user account that matches the user name and password of the administrator account that you used to log on to the Operations Center.

Create a monitoring user account

Select this option to have the wizard create an IBM Spectrum Protect Plus user account.

To enable the Operations Center to access IBM Spectrum Protect Plus and create the account, provide credentials for an IBM Spectrum Protect Plus user account that is assigned to the SYSADMIN role. Enter the credentials in the **User name** and **Password** fields as shown in the following figure.

Add Server

Authorization

Identify or create a user account on the IBM Spectrum Protect Plus server for monitoring. [Learn more](#)

Use Operations Center credentials (User account with the same credentials must already be defined on server)

Create a monitoring administrator

Specify IBM Spectrum Protect Plus login credentials for a user account that can create custom user roles and user accounts. This user account is used only during configuration. During configuration, a new user role and account for monitoring are created.

User name

Password

Figure 12. Entering IBM Spectrum Protect Plus credentials

The credentials that are entered here are not saved. The Operations Center logs on to the IBM Spectrum Protect Plus server by using these account credentials and creates the user account `OC_MONITOR_number`, where *number* is a random number for identification. The Operations Center will connect to the IBM Spectrum Protect Plus environment by using the new account.

3. Click **Add Server**.

If the operation is successful, results are displayed as shown in the following figure:

Add Server

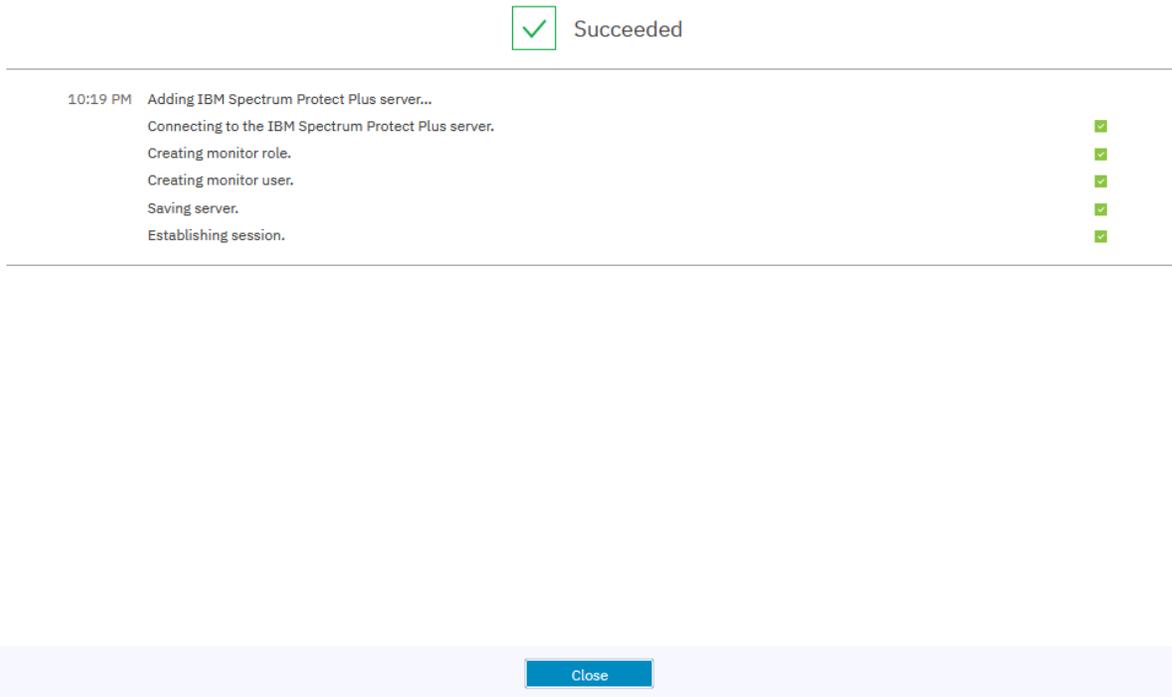


Figure 13. IBM Spectrum Protect Plus added successfully

Entering the Operations Center URL

To access the Operations Center from IBM Spectrum Protect Plus, enter the URL for the Operations Center in the IBM Spectrum Protect Plus global preferences.

About this task

You must have IBM Spectrum Protect Plus administrator credentials to configure global preferences.

When this preference is entered, the IBM Spectrum Protect icon  is active on the IBM Spectrum Protect Plus menu bar.

Procedure

To enter the URL for the Operations Center, complete the following steps:

1. In the navigation panel, click **System Configuration > Global Preferences**.
2. Enter the URL for the Operations Center in the **IBM Spectrum Protect Operations Center URL** field.

Global Preferences

Register system preferences for your IBM Spectrum Protect Plus environment.

Integration with other storage products

IBM Spectrum Protect Operations Center



<https://tapsrv09.storage.tucson.il>



URL

Figure 14. Entering the Operations Center URL

3. To activate the IBM Spectrum Protect icon on the IBM Spectrum Protect Plus menu bar, log off IBM Spectrum Protect Plus and log back on again.

Accessing the Operations Center

Start the Operations Center to monitor your IBM Spectrum Protect Plus environment.

Before you begin

Ensure that you completed the following tasks:

- [“Adding IBM Spectrum Protect Plus to the Operations Center” on page 131](#)
- [“Entering the Operations Center URL” on page 134](#)

Procedure

To access the Operations Center and monitor your IBM Spectrum Protect Plus environment, complete the following steps:

1. On the IBM Spectrum Protect Plus menu bar, click the IBM Spectrum Protect icon .
2. Log on to the Operations Center.
3. On the Operations Center menu bar, click **Overviews > Protect Plus**.

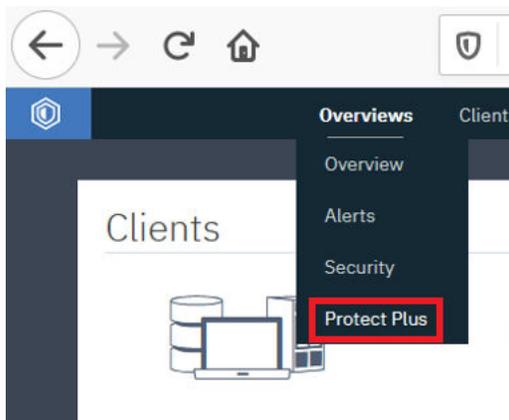


Figure 15. Selecting IBM Spectrum Protect Plus in the Operations Center

4. View the status of your IBM Spectrum Protect Plus environment on the IBM Spectrum Protect Plus monitoring dashboard as shown in the following example figure:

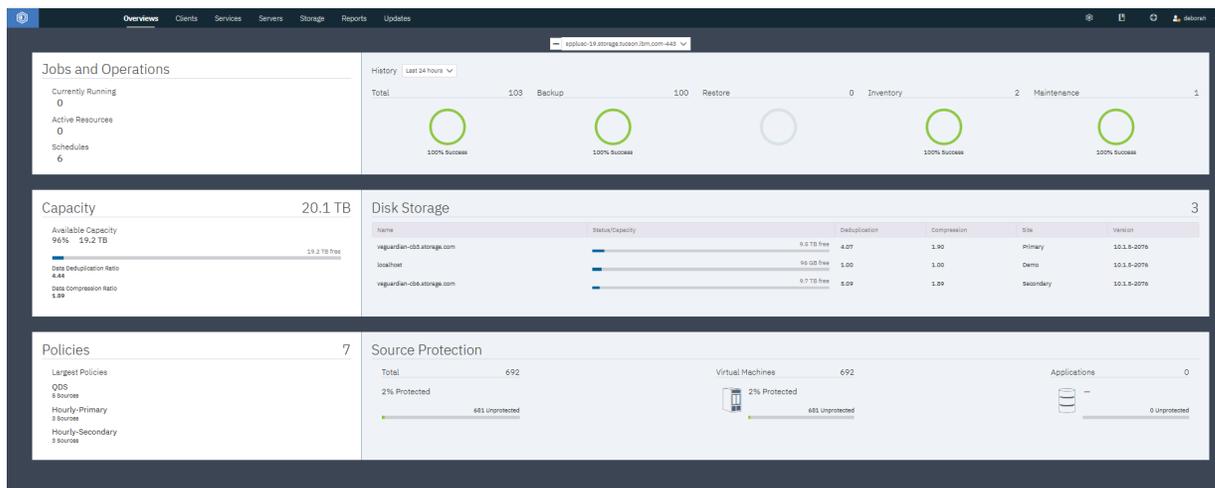


Figure 16. Viewing the IBM Spectrum Protect Plus dashboard

Managing backup storage

All IBM Spectrum Protect Plus environments must include a primary backup storage location for workload snapshots.

If your primary storage is a vSnap server, you can copy snapshots from the primary backup storage to secondary storage for longer-term data protection.

The following table shows the backup storage types that are available for IBM Spectrum Protect Plus. Depending on the workloads that you are backing up, a storage type can be available for primary backup storage only, for secondary backup storage only, or for both primary and secondary backup storage.

Backup storage type	Available as primary or secondary storage?	Description
vSnap server	Can be used as primary storage for all workloads other than container and Amazon EC2 workloads.	The vSnap server is a stand-alone appliance that is deployed virtually or installed physically on a system that meets the minimum requirements. Each vSnap server in the environment must be registered in IBM Spectrum Protect Plus. To install and manage vSnap servers, follow the instructions in Chapter 4, “Installing and managing vSnap servers,” on page 35.
Amazon Web Services (AWS) Elastic Block Store (EBS)	Available only as primary storage for Amazon EC2 workloads.	EC2 data is stored in AWS EBS snapshots. IBM Spectrum Protect Plus manages these snapshots for backup and restore operations.

Table 4. (continued)

Backup storage type	Available as primary or secondary storage?	Description
Cloud storage systems	<p>Can be used as primary storage for container workloads and the IBM Spectrum Protect Plus catalog.</p> <p>Can be used as secondary storage for data that is backed up to a vSnap server.</p>	<p>A cloud storage system is hosted by one of the following cloud providers:</p> <ul style="list-style-type: none"> • Amazon S3 • IBM Cloud Object Storage • Microsoft Azure • S3 compatible object storage providers <p>Limitations:</p> <ul style="list-style-type: none"> • For IBM Cloud Object Storage, support for retention-enabled vaults is not available. • For S3 compatible storage, generic S3 support is based on external certification processes. For the list of supported S3 compatible providers, see technote 1087149. <p>To add and manage cloud providers, see the instructions in “Managing cloud storage” on page 138.</p>
Repository server	Available only as secondary storage for data that is backed up to a vSnap server.	<p>The repository server must be IBM Spectrum Protect server 8.1.7 or later to copy data to standard object storage. To copy data to tape, IBM Spectrum Protect server 8.1.8 or later is required.</p> <p>To add and manage repository servers, see the instructions in “Managing repository server storage” on page 145.</p>
Open Snap Store Manager (OSSM)	Available only as primary storage for VMware workloads.	<p>To support direct data backup operations from IBM Spectrum Protect Plus, you must install and configure the OSSM component on the same system as the IBM Spectrum Protect server.</p> <p>To backup VMware data using Open Snap Store Manager (OSSM), see “Protecting VMware data using Open Snap Store Manager” on page 107.</p>

Adding keys and certificates for cloud and repository server backup storage providers

Before you can access a cloud or repository server storage provider, you must add the provider to IBM Spectrum Protect Plus. During the process of adding the provider, you must supply an access key to help establish a secure connection. A certificate might also be required.

You can add keys and certificates when you add the storage provider to IBM Spectrum Protect Plus or you can add them in advance.

To add keys and certificates in advance, complete the steps in the following topics:

- [“Adding an access key” on page 166](#)
- [“Adding a certificate” on page 167](#)

If you add keys and certificates in advance, select **Use existing access key** or **Use existing certificate** and select the key or certificate when you add a backup storage provider.

Related concepts

[“Copying snapshots to secondary backup storage” on page 14](#)

If your primary backup storage is a vSnap server, you can copy snapshots from the primary backup storage to secondary storage for longer-term data protection. Secondary storage is not available for container data that is backed up to cloud storage.

Managing cloud storage

You can use cloud storage as primary backup storage for container workloads and the IBM Spectrum Protect Plus catalog, or as secondary storage from the vSnap server.

The steps required to add cloud storage to IBM Spectrum Protect Plus are the same for primary and secondary storage.

Limitations:

- For IBM Cloud Object Storage, support for retention-enabled vaults is not available.
- For S3 compatible storage, generic S3 support is based on external certification processes. For the list of supported S3 compatible providers, see [technote 1087149](#).

Adding Amazon S3 Object Storage

You can add Amazon Simple Storage Service (S3) as primary backup storage for container workloads and the IBM Spectrum Protect Plus catalog, or as secondary storage from the vSnap server.

Before you begin

Configure the key that is required for the cloud object. For instructions, see [“Adding an access key” on page 166](#).

Ensure that cloud storage buckets are created for the IBM Spectrum Protect Plus data. For instructions about creating buckets, see [Amazon Simple Storage Service Documentation](#).

Procedure

To add Amazon S3 cloud storage as a backup object storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Cloud storage**.
2. Click **Add cloud storage** to open the **Add cloud storage** wizard.
3. Click **Amazon S3**, and then click **Next**.
4. Complete the fields on the **Cloud details** page, and then click **Next**:

Name

Enter a meaningful name that helps to identify the cloud storage.

Use existing access key

Enable this option to select a previously entered key for the storage, and then select the key from the **Select a key** list.

If you do not select this option, complete the following fields to add a key:

Key name

Enter a meaningful name to identify the key.

Access key

Enter the AWS access key. Access keys are created in the AWS Management Console.

Secret key

Enter the AWS secret key. Secret keys are created in the AWS Management Console.

5. Complete the fields on the **Get buckets** page, and then click **Next**:

Region

Select the region for the cloud storage, and then click **Update buckets** to select the buckets that you want to use for storage backup, copy, and archive operations. The buckets that you select depend on the backup configuration that you want to use.

If you are backing up container workloads, you can use cloud storage as your primary backup storage.

If your primary backup storage location is a vSnap server, you can use cloud storage as a copy or archive location.

Backup object storage bucket

Select a bucket to serve as the backup storage target.

Standard object storage bucket

Select a bucket to serve as the copy target.

Archive object storage bucket

Select a cloud storage resource to serve as the archive target.

Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copying snapshots to secondary backup storage”](#) on page 14.

Deep archive

Select to register Amazon S3 Glacier Deep Archive buckets for long-term archiving. This field is optional.

6. Review your selections, and then click **Submit**.

The cloud storage is added to the cloud servers table.

What to do next

After you add the S3 storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see “Creating an SLA policy for databases and file systems” on page 206. To modify an existing SLA policy, see “Editing an SLA policy” on page 210.

Adding IBM Cloud Object Storage as a backup storage provider

You can add IBM Cloud Object Storage as primary backup storage for container workloads and the IBM Spectrum Protect Plus catalog, or as secondary storage from the vSnap server.

Before you begin

Configure the key and certificate that are required for the cloud object. For instructions, see [“Adding an access key” on page 166](#) and [“Adding a certificate” on page 167](#).

Ensure that there are cloud storage buckets created for the IBM Spectrum Protect Plus data before you add the cloud storage in the following steps. For information how to create buckets, see [About IBM Cloud Object Storage](#).

When creating a bucket on IBM Cloud Object Storage (COS), ensure that both **Add Archive rule** and **Add Expiration rules** are not selected when creating buckets that are to be used for copy or archive. This can result in a failure with the “bucket has an unsupported lifecycle configuration” error when the job attempts to run in IBM Spectrum Protect Plus. The **Add Retention policy** option may be set for a bucket to be used for copy, but should not be set for a bucket that will be used for archiving.

The Cold Vault bucket of type should only be used when archiving, as it is the lowest-cost option and is described as ideal for long-term retention of data that will be minimally accessed.

For on-premises IBM Cloud Object Storage (COS), the user associated with the access key and secret key that is registered in IBM Spectrum Protect Plus must be assigned as an Owner of the vault in the IBM COS Manager interface. It is not sufficient for the user to have only Read/Write access to the vault.

When adding IBM Cloud Object Storage (COS), the method for obtaining the access and secret key will depend on the deployment model. If on-premise, keys can be obtained from the IBM COS Manager Console. For IBM COS IaaS, keys are created when a service account is created and can be obtained from the softlayer portal. If using IBM COS (COS as a Service), the access and secret key are not created by default; when a service account is created, check the **Include HMAC Credential** box, and add `{“HMAC”: true}` to the **Add Inline Configuration Parameters** text area.

Procedure

To add IBM Cloud Object Storage as a backup storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Cloud storage**.
2. Click **Add cloud storage** to open the **Add cloud storage** wizard.
3. Click **IBM Cloud Object Storage**, and then click **Next**.
4. Complete the fields on the **Cloud details** page, and then click **Next**:

Name

Enter a meaningful name to help identify the cloud storage.

Use existing access key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

If you do not select this option, complete the following fields to add a key:

Key name

Enter a meaningful name to identify the key.

Access key

Enter the access key.

Secret key

Enter the secret key.

Certificate

Select a method of associating a certificate with the resource:

Upload

Select and click **Browse** to locate the certificate, and then click **Upload**.

Copy and paste

Select to enter the name of the certificate, copy and paste the contents of the certificate, and then click **Create**.

Use existing

Select to use a previously uploaded certificate.

A certificate is not required if you are adding public IBM Cloud Object Storage.

5. Complete the fields on the **Get buckets** page, and then click **Next**:

Endpoint

Enter the endpoint path for the cloud storage, and then click **Update buckets** to select the buckets that you want to use for storage backup, copy, and archive operations. The buckets that you select depend on the backup configuration that you want to use.

If you are backing up container workloads, you can use cloud storage as your primary backup storage.

If your primary backup storage location is a vSnap server, you can use cloud storage as a copy or archive location.

Backup object storage bucket

Select a bucket to serve as the backup storage target.

Standard object storage bucket

Select a bucket to serve as the copy target.

Archive object storage bucket

Select a cloud storage resource to serve as the archive target.

Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copying snapshots to secondary backup storage” on page 14](#).

6. Review your selections, and then click **Submit**.

The cloud storage is added to the cloud servers table.

What to do next

After you add the IBM Cloud Object Storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see “Creating an SLA policy for databases and file systems” on page 206 . To modify an existing SLA policy, see “Editing an SLA policy” on page 210 .

Adding Microsoft Azure cloud storage as a backup storage provider

You can add Microsoft Azure cloud storage as primary backup storage for container workloads and the IBM Spectrum Protect Plus catalog, or as secondary storage from the vSnap server.

Before you begin

Ensure that there are cloud storage buckets created for the IBM Spectrum Protect Plus data before you add the cloud storage in the following steps. For information about how to create buckets, see the Azure documentation.

Procedure

To add Microsoft Azure cloud storage as a backup storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Cloud storage**.
2. Click **Add cloud storage** to open the **Add cloud storage** wizard.
3. Click **Microsoft Azure Blob Storage**, and then click **Next**.
4. Complete the fields on the **Cloud details** page, and then click **Next**:

Name

Enter a meaningful name to identify the cloud storage.

Use existing access key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

If you do not select this option, complete the following fields to add a key:

Key name

Enter a meaningful name to identify the key.

Storage account name

Enter the Microsoft Azure access storage account name. This name can be obtained from the Azure Management Portal.

Storage account shared key

Enter the Microsoft Azure key from any one of the key fields in the Azure Management Portal: key1 or key2.

5. Complete the fields on the **Get buckets** page, and then click **Next**:

Endpoint

Select the endpoint for the cloud storage, and then click **Update buckets** to select the buckets that you want to use for storage backup, copy, and archive operations. The buckets that you select depend on the backup configuration that you want to use.

If you are backing up container workloads, you can use cloud storage as your primary backup storage.

If your primary backup storage location is a vSnap server, you can use cloud storage as a copy or archive location.

Backup object storage bucket

Select a bucket to serve as the backup storage target.

Standard object storage bucket

Select a bucket to serve as the copy target.

Archive object storage bucket

Select a cloud storage resource to serve as the archive target.

Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copying snapshots to secondary backup storage”](#) on page 14.

6. Review your selections, and then click **Submit**.

The cloud storage is added to the cloud servers table.

What to do next

After you add the Microsoft Azure storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	<p>To create an SLA policy, see “Creating an SLA policy for databases and file systems” on page 206.</p> <p>To modify an existing SLA policy, see “Editing an SLA policy” on page 210.</p>

Adding S3 compatible object storage

You can add the S3 compatible storage providers that are listed in [technote 1087149](#). Support for S3 compatible providers is based on external certification processes.

Before you begin

Configure the key that is required for the cloud object. For instructions, see [“Adding an access key”](#) on page 166.

Ensure that cloud storage buckets are available. For more information about cloud storage buckets, see the documentation for the S3 compatible storage provider.

Procedure

To add S3 compatible cloud storage as a backup target, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Cloud storage**.
2. Click **Add cloud storage**.
3. Click **S3 Compatible Storage**, and then click **Next**.
4. Complete the fields on the **Cloud details** page, and then click **Next**:

Name

Enter a meaningful name to identify the cloud storage.

Use existing access key

Enable this option to select a previously entered key for the storage, and then select the key from the **Select a key** list.

If you do not select this option, complete the following fields to add a key:

Key name

Enter a meaningful name to identify the key.

Access key

Enter the S3 compatible access key. For instructions about obtaining access keys, see the documentation for the S3 compatible storage provider.

Secret key

Enter the S3 compatible secret key. For instructions about obtaining access keys, see the documentation for the S3 compatible storage provider.

Certificate

Select the appropriate option to add a certificate for the S3 compatible storage:

Upload

To upload a certificate, click **Browse** to locate and select the certificate. Click **Upload**.

Copy and paste

Enter a name for the certificate and paste the certificate into the text area. Click **Create**.

Use existing

If a certificate exists, select the certificate from the **Select a certificate** list.

5. Complete the fields on the **Get buckets** page, and then click **Next**:

Endpoint

Enter the endpoint path for the cloud storage, and then click **Update buckets** to select the buckets that you want to use for storage backup, copy, and archive operations. The buckets that you select depend on the backup configuration that you want to use.

If you are backing up container workloads, you can use cloud storage as your primary backup storage.

If your primary backup storage location is a vSnap server, you can use cloud storage as a copy or archive location.

Backup object storage bucket

Select a bucket to serve as the backup storage target.

Standard object storage bucket

Select a bucket to serve as the copy target.

Archive object storage bucket

Select a cloud storage resource to serve as the archive target.

Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copying snapshots to secondary backup storage”](#) on page 14.

6. Review your selections, and then click **Submit**.

The cloud storage is added to the cloud servers table.

What to do next

After you add the S3 compatible storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see “Creating an SLA policy for databases and file systems” on page 206. To modify an existing SLA policy, see “Editing an SLA policy” on page 210.

Editing settings for cloud storage

Edit the settings for a cloud storage provider to reflect changes in your cloud environment.

Procedure

To edit a cloud storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Cloud storage**.
2. Select the cloud storage provider, and then click **Edit** to open the **Edit cloud storage** wizard.
3. Revise the settings for the cloud provider, and then click **Save**.

Deleting cloud storage

Delete a cloud storage provider to reflect changes in your cloud environment. Ensure that the provider is not associated with any SLA policies before deleting the provider.

Before you begin



Attention: Deleting a cloud storage provider could result in a loss of data. Ensure that any required data is backed up before you delete the cloud storage provider.

Procedure

To delete a cloud storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Cloud storage**.
2. Select the cloud storage provider, and then click **Remove**.
3. Click **Yes** to delete the provider.

Managing repository server storage

You can copy data to a repository server for longer-term data protection. The repository server must be IBM Spectrum Protect server 8.1.7 or later to copy data to standard object storage. To copy data to tape, IBM Spectrum Protect server 8.1.8 or later is required.

You can choose to replicate the IBM Spectrum Protect Plus data that is copied to the IBM Spectrum Protect server to a target server. However, IBM Spectrum Protect Plus is not aware of subsequent IBM Spectrum Protect server replication operations and you cannot restore the replicated data from the target IBM Spectrum Protect server to IBM Spectrum Protect Plus.

Configuration for copying or archiving data to IBM Spectrum Protect

If you are planning to copy or archive IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, there are three possible configurations. Choosing which one to configure depends on which scenario applies to your data protection needs. For each scenario, there are steps that are required in both the IBM Spectrum Protect Plus and IBM Spectrum Protect server environments to complete the setup.

Tasks for configuring IBM Spectrum Protect

You must configure the IBM Spectrum Protect server to communicate with the IBM Spectrum Protect Plus server, and to enable process requests for backup and restore operations. The Amazon Simple Storage Service (S3) protocol enables communication between the two servers.

User scenario	Purpose	Steps
Copying to standard object storage when you are running daily or less frequent copies to standard object storage.	Copy data to standard object storage. In the first copy operation, a full backup copy is created. Subsequent copies are incremental. Copying data to standard object storage is useful if you want relatively fast backup and recovery times and do not require the longer-term protection, cost, and security benefits that are provided by tape storage.	To copy data to standard object storage to the IBM Spectrum Protect server, you must create a cloud-container or directory-container storage pool, and set up the object agent component of IBM Spectrum Protect. Adding the object agent is a mandatory step. In addition to setting up the required storage pool, follow steps 2-4 listed, here .

User scenario	Purpose	Steps
<p>Copying to tape when you are creating a weekly or less frequent full-copy of your data to tape storage.</p> <p>Important: Archiving data to tape cannot be run more frequently than once a week. Recovery time objectives (RTO) should be considered when recovering data from archive copies in your disaster recovery action plan. Therefore, for disaster recovery, recovering from archive data should only be used as a last resort.</p>	<p>When you copy data to tape, a full copy of the data is created at the time of the copy process. Copying data to tape provides extra security benefits. By storing tape volumes at a secure, offsite location that is not connected to the internet, you can help to protect your data from online threats such as malware and hackers. However, because copying to these storage types requires a full data copy, the time that is required to copy data increases. In addition, the recovery time can be unpredictable and the data might take longer to process before it is usable. Some of the data may be duplicated on the tape storage pool and cache storage pool.</p>	<p>To copy data to tape, you must create a tape storage pool first and then you must create a disk storage pool which is where the cold-data-cache storage pool will reside on the IBM Spectrum Protect server. Adding the object agent is a mandatory step. Follow steps 1-4 listed, here.</p>
<p>Mixture of both standard object storage and long-term copying to tape</p>	<p>Secure your data in incremental backups on the IBM Spectrum Protect server, as well as retaining data on tape for longer term security.</p>	<p>This is a combination of the previous cases: data is stored to tape and data is stored on standard object storage at the IBM Spectrum Protect server. As well as setting up the required data storage pools for both scenarios, the creation of an object agent is mandatory.</p>

The four steps required to set up and configure the data transfer communication between IBM Spectrum Protect Plus and the IBM Spectrum Protect server are as follows:

1. If you are setting up storage pools for copying data to tape follow Step1. Create storage pools on the IBM Spectrum Protect server by using the IBM Spectrum Protect Operations Center. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 147](#). This step is required only if you are setting IBM Spectrum Protect for archiving with copies run once a week or less frequently.
2. Create a policy domain that points to the storage pool or pools. The policy domain defines the rules that control the backup services for IBM Spectrum Protect Plus. For instructions, see [“Step 2: Configuring an object policy domain” on page 148](#).
3. If you are copying data to a standard storage pool or to tape, you must add standard object storage on the IBM Spectrum Protect server. For instructions, see [“Step 3: Setting up standard object storage” on page 150](#).
4. Add an object agent on the IBM Spectrum Protect server. The object agent provides a gateway between the IBM Spectrum Protect Plus server and the IBM Spectrum Protect server. For instructions, see [“Step 4: Adding an object agent for copying data ” on page 153](#).
5. To complete the setup, you must add an object client on the IBM Spectrum Protect server. The object client identifies the IBM Spectrum Protect Plus server and enables it to store objects at the IBM Spectrum Protect server. The same credentials as those that you used for IBM Spectrum Protect Plus are used for the object client, which is the object client that is associated with the policy domain as set

up in Step 2. For instructions to set up an object client, see [“Step 5: Adding and configuring an object client for copying data”](#) on page 155.

Tip: Alternatively, enter the **DEFINE STGPOOL** command to create a storage pool as described in the following topics:

What to do next

1. After you complete the tasks required for IBM Spectrum Protect storage, you must add the IBM Spectrum Protect server to IBM Spectrum Protect Plus. For information about how to do this, follow the instructions in [“Registering a repository server as a backup storage provider”](#) on page 156.
2. When that is done, you can create an SLA policy that defines the IBM Spectrum Protect server as the backup storage target. For more information to help you choose which type of policy you need, see [“Configuration for copying or archiving data to IBM Spectrum Protect”](#) on page 145

Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape

Before you can copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server for archiving purposes, you must configure an object agent service. For long-term archiving of data, you must configure a cold data storage pool. If you are not planning to archive data to tape on the IBM Spectrum Protect server, you can skip this step.

About this task

Before you start, ensure that you have sized your cold cache storage needs by using the sizing tool and the Blueprints. For information about how to do this, see the [Blueprints](#). For more useful links and videos, see [“Deployment storyboard for IBM Spectrum Protect Plus”](#) on page 1.

Object client data that is specified with an S3 Glacier storage class is not frequently accessed. To enable the copying of this data, which is often called *cold data*, to tape storage, the data is written temporarily to a storage pool that meets the requirements for handling object data. The data is then moved to the tape device or VTL. This storage pool, called a *cold-data-cache storage pool*, is assigned to a policy domain for object clients. Only data from object clients can be written to or restored from a cold-data-cache storage pool.

Procedure

If you are not using the Operations Center, you can use the **define stgpool** command. The command can be defined as follows:

```
define stgpool NAME
stgtype=colddatacache
```

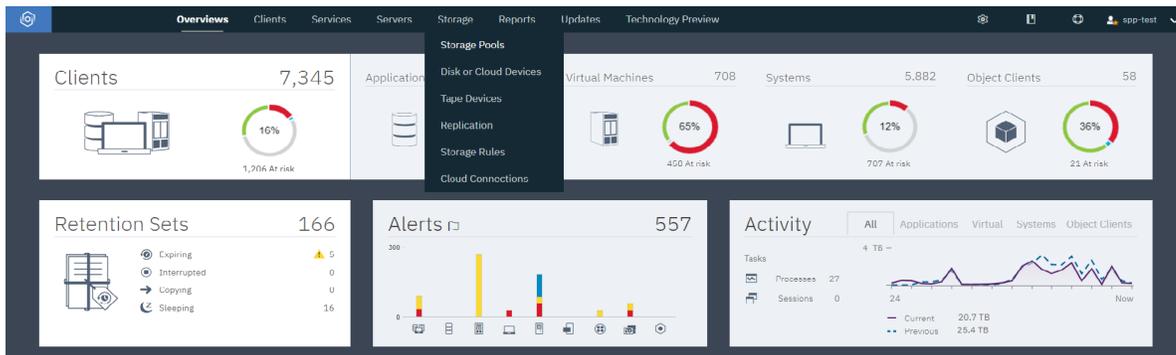
Note: To configure standard pools for object storage, follow these steps but when you define the type of storage pool, select Standard.

To configure the IBM Spectrum Protect server to copy data from an object client to physical tape media or a VTL, complete the following configuration steps:

1. On the IBM Spectrum Protect server, configure a primary storage pool that represents a tape device or VTL. This primary storage pool is the destination for the object data that you want to copy.
Later, when you define the cold-data-cache storage pool, you must specify this tape pool as the next storage pool for the cold-data-cache pool.

Restrictions: The following restrictions apply to the tape storage pool:

- You cannot replicate object client data to or from the tape storage pool.
 - The tape storage pool cannot be deduplicated.
 - A next storage pool cannot be specified for the tape storage pool.
- a) On the Operations Center menu bar, click **Storage > Storage Pools**.



- b) On the **Storage Pools** page, click **Storage Pool**.
 - c) In the **Add Storage Pool** wizard, select **Object Client** to enable object clients to copy data to tape.
2. Step through the wizard steps to configure a cold-data-cache storage pool.

A cold-data-cache storage pool consists of one or more file system directories on disk. It is an intermediary storage pool between the object client and a tape device or VTL and is linked to the primary sequential access storage pool that represents the tape device or VTL. Identify one or more existing file system directories for temporary disk storage and the primary sequential access storage pool that represents the tape device or VTL.

3. On the **Cold Data Cache** page, specify one or more existing file system directories for disk storage. Enter a fully qualified path name that conforms to the syntax that is used by the server operating system.

For example, enter `c:\temp\dir1\` for Microsoft Windows, or `/tmp/dir1/` for UNIX.

The object data is stored in sequential volumes in the file system directories. An object client can copy infrequently accessed data, or cold data, to physical tape media or to a VTL. When an object client copies cold data, the data is first stored in the cold data cache. The data is then migrated, without a migration delay, to the primary tape storage pool that represents the physical tape media or VTL. After the data is migrated to tape, it is deleted from the cold data cache. The cold data cache is used as a staging area for restoring cold data to the object client. During restore operations, the data is copied to the cold data cache. The data remains in the cold data cache for a period that is specified by the object client. Data is restored to the object client from the cold data cache, and not directly from the tape or VTL.

If you specify multiple directories for performance enhancement, ensure that the directories correspond to separate physical volumes. Although the cold data cache is used for temporary storage, it must be large enough to hold the data that is copied from the object client before the data is migrated to tape. It must also be large enough to hold data during restore operations for the period that is specified by the object client.

What to do next

When you complete the configuration of the cold data cache storage pool, create the object domain. For instructions about how to do that, see [“Step 2: Configuring an object policy domain”](#) on page 148.

Step 2: Configuring an object policy domain

Before you copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, you must create and configure an object policy domain. The policy domain defines the rules that control the backup services for IBM Spectrum Protect Plus. You must add a standard storage pool which is with a directory or cloud container based storage for copies, and a cold pool if you are copying data to tape or archiving data.

Procedure

1. Verify the settings for the policy domain that you plan to use for copying data. Object clients that are defined or updated in the IBM Spectrum Protect server 8.1.8 or later must be assigned to policy domains that are created with the **DEFINE OBJECTDOMAIN** command. An object client node is

associated with this policy domain when the node is registered or updated with the **REGISTER NODE** or **UPDATE NODE** command.

Restriction: Beginning with IBM Spectrum Protect server 8.1.8, all new object client nodes must be assigned to object policy domains.

For object client nodes that were assigned to non-object policy domains before v8.1.8, you do not have to update the assignment after you upgrade the server to IBM Spectrum Protect server 8.1.8. However, if any update to the object client node's domain is required, the node must be assigned to an object policy domain.

2. Review the following considerations for specifying policy domains for copy operations.

- For IBM Spectrum Protect server, a policy domain can specify management classes for standard storage pools (cloud-container or directory-container storage pools), cold-data-cache storage pools, or both standard and cold-data-cache storage pools.

However, to copy data from IBM Spectrum Protect Plus, you must specify the following management classes depending on whether you are copying data to a cloud-container or directory-container storage pool or are copying data to a cold-data-cache storage pool for storage on physical tape media or in a virtual tape library (VTL):

- To copy data to a cloud-container or directory-container storage pool, use the **STANDARDPOOL** parameter to define the storage pool for the policy domain as shown in the following example:

```
define objectdomain mydomain standardpool=hotpool
```

- To copy data to a cold-data-cache storage pool, you must specify both a standard pool and a cold pool for the policy domain. A standard pool is required to store metadata that is used for restore and other IBM Spectrum Protect Plus operations. To define a cold-data-cache storage pool for the policy domain, use the **COLDPOOL** parameter, as shown in the following example:

```
define objectdomain mydomain standardpool=hotpool coldpool=coldpool
```

- All objects are uniquely named. There are no inactive versions of objects. When you define a policy domain, the following Storage Management policies are specified automatically:
 - The `Versions Data Exists` field is set to 1.
 - The `Retain Extra Versions` and the `Retain Only Version` fields are set to 0.
- The IBM Spectrum Protect Plus server controls the time when objects are deleted.

Example: Display detailed information about a policy domain for an IBM Spectrum Protect Plus copy operation

When the policy domain was created, it was assigned management classes and copy groups. You can use the **QUERY COPYGROUP** command to view information about the destination storage pools for the policy domain. In the following example, the policy domain name is XYZ. The destination storage pools are HOTPOOL and COLDPOOL.

```
query copygroup xyz standard f=d
```

```

Policy Domain Name: XYZ
Policy Set Name: STANDARD
Mgmt Class Name: COLD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 1
Versions Data Deleted: 1
Retain Extra Versions: 0
Retain Only Version: 0
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: COLDPPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/22/20 17:03:46
Managing profile:
Changes Pending: No

Policy Domain Name: XYZ
Policy Set Name: STANDARD
Mgmt Class Name: STANDARD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 1
Versions Data Deleted: 1
Retain Extra Versions: 0
Retain Only Version: 0
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: HOTPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 03/05/20 22:15:18
Managing profile:
Changes Pending: No

```

What to do next

After you create the object domain, proceed to the next step [“Step 3: Setting up standard object storage”](#) on page 150.

Step 3: Setting up standard object storage

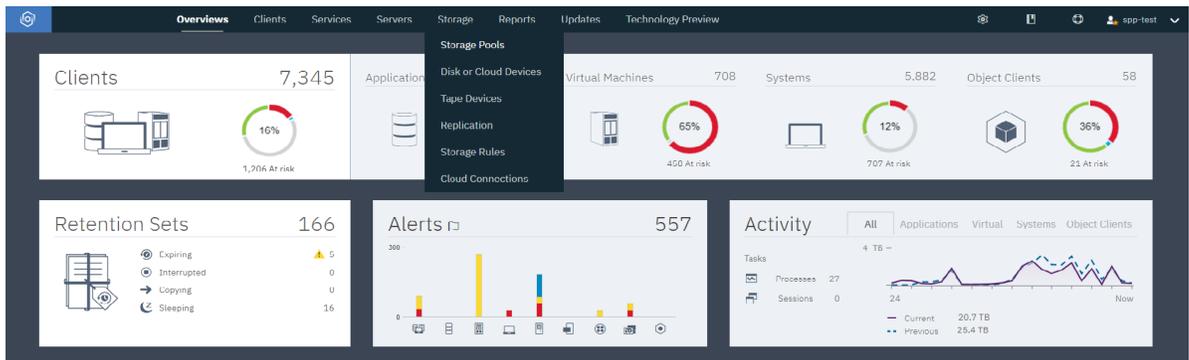
To set up standard object storage for copying data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, log in to the Operations Center and follow the procedure to set up storage pools. Complete the process by following the steps to create an object agent service by using the Operations Center wizard.

Before you begin

Before you start you must set up storage pools for standard storage or for copying to tape. If you are copying to tape, you must set up the cold data cache storage pool, and for standard object storage you must create and configure storage pools as required. For instructions about how to set up the cold data cache storage pool, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape”](#) on page 147.

Procedure

1. Create a directory-container storage pool by completing the following steps:
 - a) On the Operations Center menu bar, click **Storage > Storage Pools**.



b) On the **Storage Pools** page, click **Storage Pool**.

c) Complete the steps in the **Add Storage Pool** wizard.

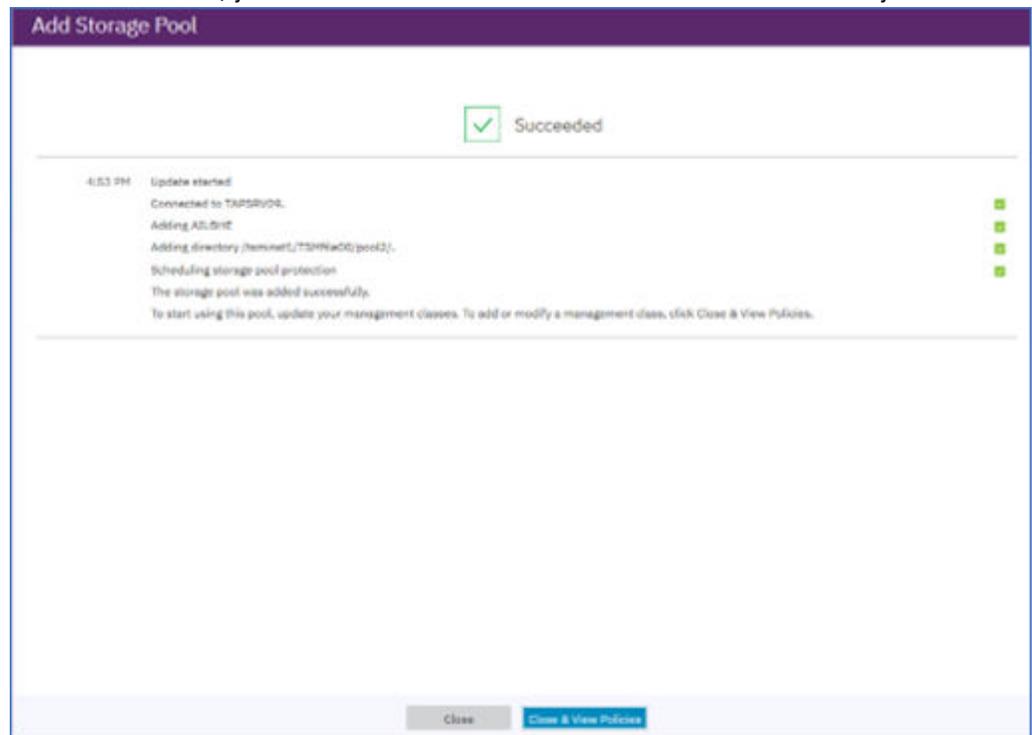
Tip: Select **Directory** for the type of container-based storage, and add directories with the + icon. Click **Next** to continue.

d) Review the **Protect Pool** summary, and click **Next**.

e) Specify an overflow pool is that is required.

f) Click **Add Storage Pool** to complete the creation of the storage pool.

If the operation was successful, you will see an icon to indicate success with a summary of the

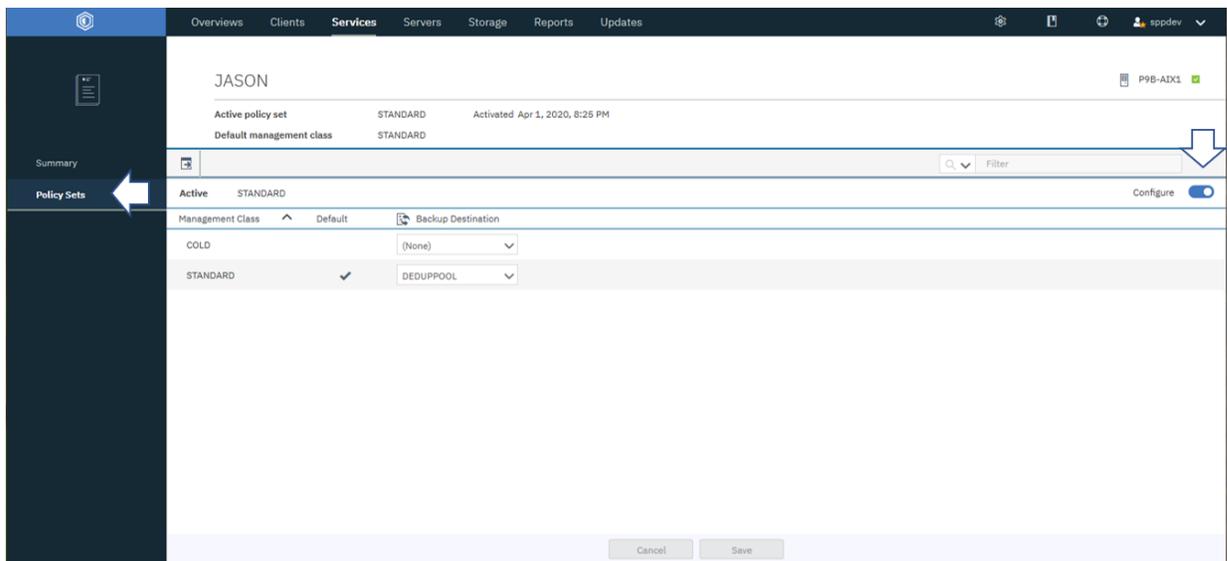


storage pool.

2. In the **Services > Policies** page, select a policy, and click **Details**.

Policy Domain	Server	Clients	Mgmt Classes	Option Sets	Schedules	Default Mgmt Class	Backup Destination	Archive Destination
207695	ION	1	1	0	0	207695	207695	207695
ADP	PROTO	1	4	0	0	LT05POOL	LT05POOL	
APITESTDOM	PROTO	3	6	1	0	NOTHING1	BACKUPPOOL	ARCHIVEPOOL

- You can edit an existing domain policy by following these steps:
 - Update one or more management classes to use the new pool by editing the **Backup Destination** field of the table.
 - Click **Save**.
 - Or, you can create a new domain by running the **define objectdomain** command. For more information, see the previous step [“Step 2: Configuring an object policy domain”](#) on page 148.
3. On the **Details** page, click **Policy Sets**. Click the **Configure** toggle to make the policy sets editable.



4. Change the Backup Destination to the newly created storage pool, or add a new management class,



- Click **Activate**.
Changing the active policy set might result in data loss. A summary of the differences between the active policy set and the new policy set is displayed before the change is made.
- Review the differences between corresponding management classes in the two policy sets, and consider the consequences on client files. Client files that are bound to management classes in the currently active policy set are, after activation, bound to the management classes with the same names in the new policy set.
- Identify management classes in the currently active policy set that do not have counterparts in the new policy set, and consider the consequences on client files. Client files that are bound to these management classes are, after activation, managed by the default management class in the new policy set.
- If the changes implemented by the policy set are acceptable, select the **I understand that these updates can cause data loss** checkbox and click **Activate**.

What to do next

Create and configure an object client for the storage pool or pools you created. For more information, see [“Step 5: Adding and configuring an object client for copying data” on page 155](#)

Step 4: Adding an object agent for copying data

Before you can copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, you must add and configure the object agent. This step is the fourth step in setting up IBM Spectrum Protect Plus with the IBM Spectrum Protect server for archiving data or copying data to object storage.

Before you begin

Ensure that the following steps are complete before you start to create the object client.

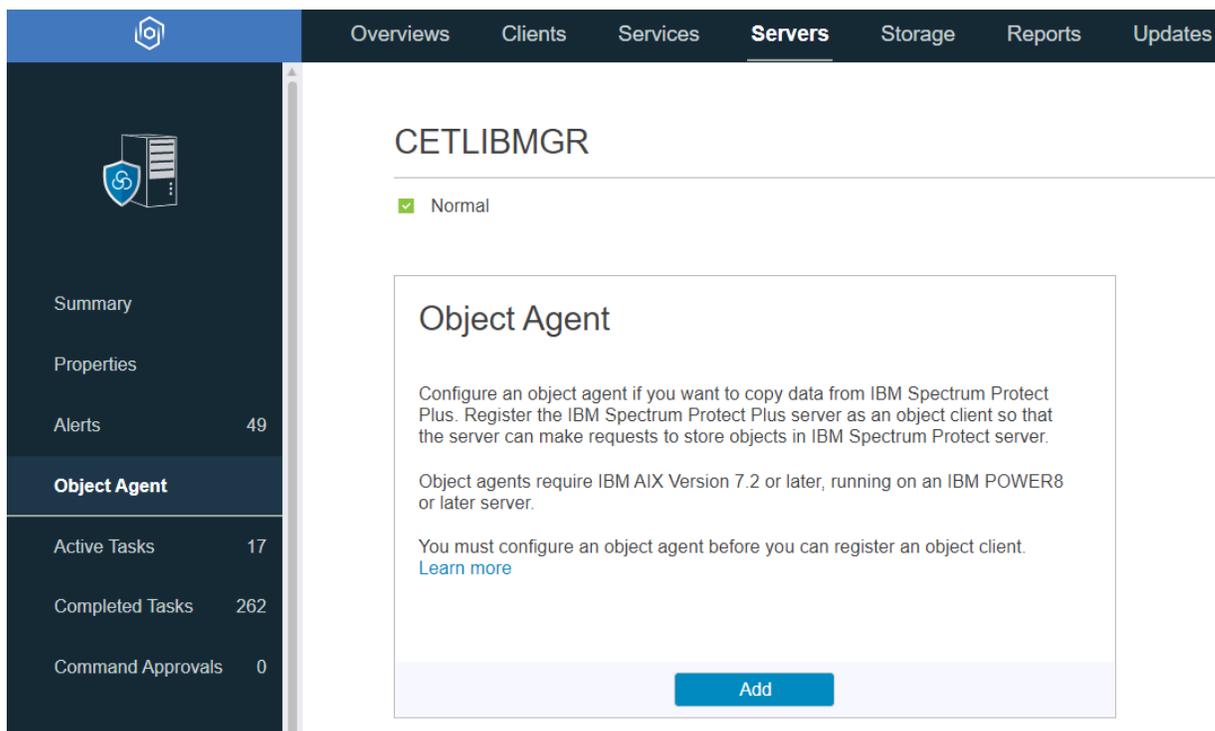
1. Ensure that you are logged in to the IBM Spectrum Protect server with an instance user ID.
2. Ensure that you have set up storage pools either for standard storage or for copying to tape. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 147](#) or [“Step 3: Setting up standard object storage” on page 150](#).
3. Ensure that you have created an object domain.

About this task

This procedure is based on an environment where the IBM Spectrum Protect server is installed on an IBM AIX® operating system AIX Version 7.2 TL 1 and SP 4 or later, running on an IBM POWER8® or later server. (LINK TO a previous version)

Procedure

1. On the Operations Center menu bar, click **Servers** .
2. Select a server and click **Details**.
3. From the navigation panel, click **Object Agent**; click **Add** to add an object agent.

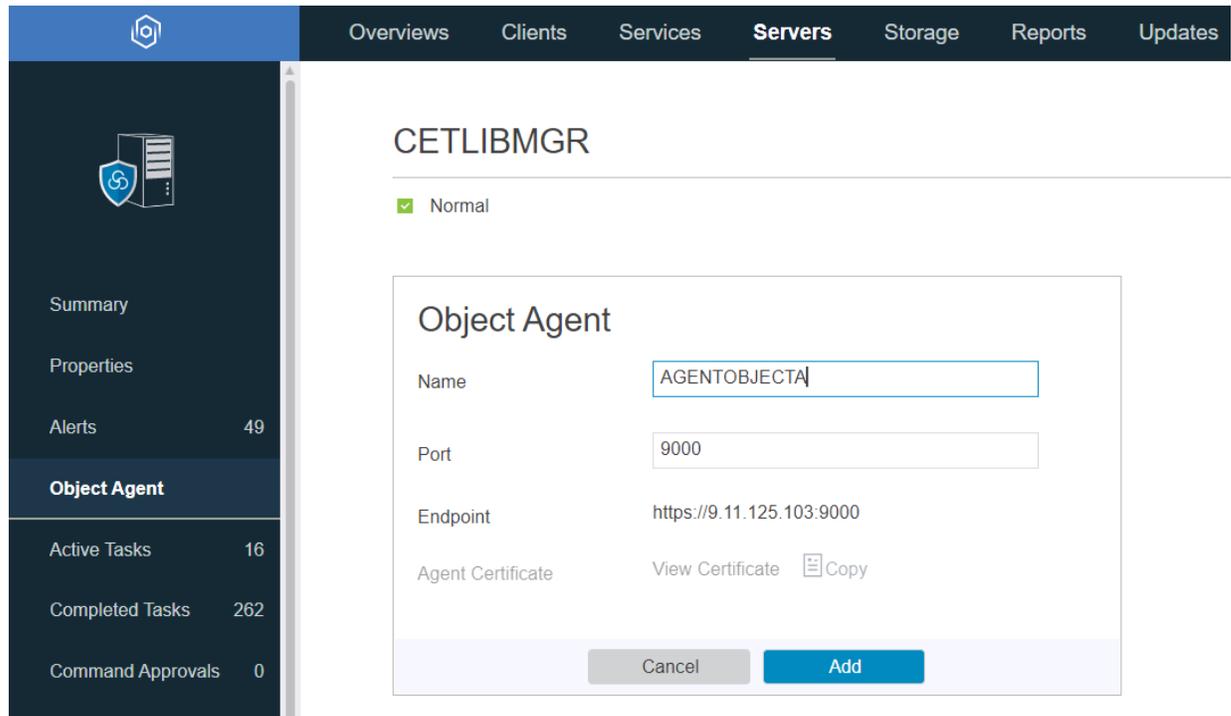


The screenshot shows the IBM Spectrum Protect Plus web interface. The top navigation bar includes 'Overviews', 'Clients', 'Services', 'Servers', 'Storage', 'Reports', and 'Updates'. The 'Servers' tab is active, and the server 'CETLIBMGR' is selected. The left navigation panel shows 'Object Agent' as the active section. The main content area displays the 'Object Agent' configuration page for 'CETLIBMGR', which is in a 'Normal' state. The page includes instructions on how to configure an object agent and a blue 'Add' button at the bottom.

Tip: If you are using the command line, run the **DEFINE SERVER** command to create an object agent. Specify OBJECTAGENT=YES. Follow the instructions in the command output. When these actions

are completed, the object agent service automatically starts on the system that is hosting the IBM Spectrum Protect server.

4. To authenticate to the object agent, use the certificate that is generated.



5. Install the object agent service by running the command that can be copied from the wizard like in the following examples:

```
[root@servername-os: /]# /opt/tivoli/tsm/server/bin/spObjectAgent service install
/home/tsminst1/tsminst1/SPPOBJAGENT/spObjectAgent_SPPOBJAGENT_1500.config
2020-03-31 15:50:07.631021 I | Installed and started system service as
nameportnumberobjectagentname
```

Here is an example

```
[root@p9b-aix1: /]# /opt/tivoli/tsm/server/bin/spObjectAgent service install
/home/tsminst1/tsminst1/SPPOBJAGENT/spObjectAgent_SPPOBJAGENT_1500.config
2020-03-31 15:50:07.631021 I | Installed and started system service as spoa9000SPPOBJAGENT
```

6. Complete the configuration by starting an object agent service by running the **startObjectAgent** command. Here is an example for **AGENTOBJECTA** object agent.

```
"/opt/tivoli/tsm/server/bin/spObjectAgent" service install
"/home/tsminst1/tsminst1/AGENTOBJECTA/spobjectAgent_AGENTOBJECTA_1500.config"
```

7. Set up the object agent service to start automatically on startup by running a command similar to the following command for AIX:

```
spobj:2:once:/usr/bin/startsrc -s nameportnumberobjectagentname
```

Here is an example:

```
spobj:2:once:/usr/bin/startsrc -s spoa9000SPPOBJAGENT
```

Step 5: Adding and configuring an object client for copying data

Before you can copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, you must configure the object client. This step is the last step in setting up the IBM Spectrum Protect server for archiving and copying of data with the Operations Center.

Before you begin

Ensure that the following steps are complete before you start to create the object client.

1. Ensure that you are logged in to the IBM Spectrum Protect server with an instance user ID.
2. Ensure that the storage pools for either standard storage or for copying to tape are set up and ready. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape”](#) on page 147 or [“Step 3: Setting up standard object storage”](#) on page 150.
3. Ensure that an object domain and an object agent are created before you start.

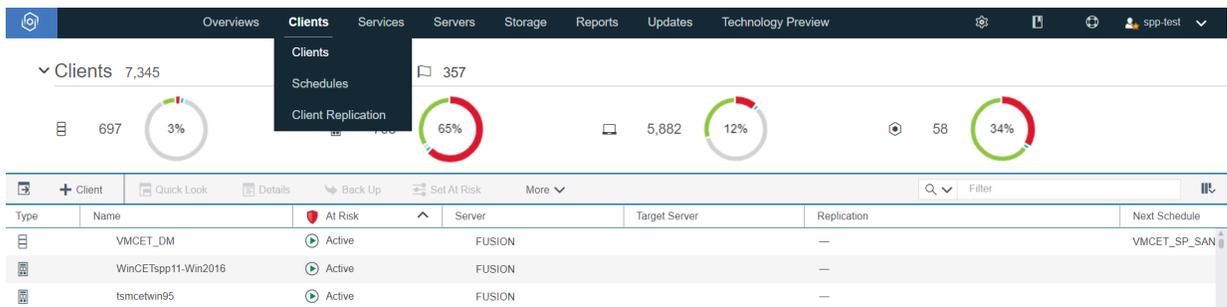
Tip: If you create an object client before you create the corresponding object agent, the **Add Client** wizard forces the creation of the object agent.

About this task

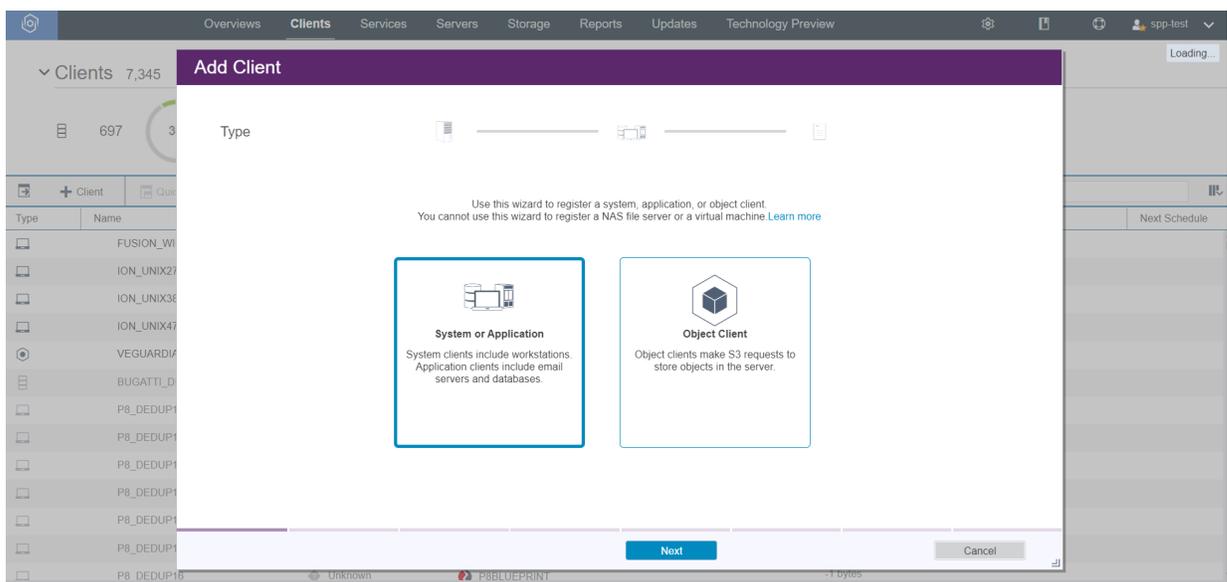
This procedure is based on an environment where the IBM Spectrum Protect server is installed on an IBM AIX operating system AIX Version 7.2 TL 1 and SP 4 or later, running on an IBM POWER8 or later server.

Procedure

1. On the Operations Center menu bar, click **Clients**.



2. Click **Client** to add a client as shown.



3. Select **Object Client** and click **Next** to start the **Add Client** wizard.

In the wizard screens, you are asked for to make the following choices and definitions for the client you are setting up.

- You can also choose to enable replication for this client.
- You must assign a client name and contact name, and an email address for reporting which you define in the final step of the wizard.
- You must assign a policy domain, which you set up in step 2, [“Step 2: Configuring an object policy domain”](#) on page 148.
- You can define at risk reporting for the client, such as a once-a-day report to the email address that you specified.

4. Click **Add Client**.

Note:

After the process finishes, you are provided with the endpoint for communicating with the object agent on the server, the access key ID, the secret access key, and the certificate for connecting securely. When IBM Spectrum Protect Plus is an object client, it directs requests to the endpoint, and uses this information in the form of the access key ID, the secret access key, and the secure certificate.

Important: Ensure that a copy of each credential is saved to a secure location.

Tip: If you are using the command line, run the **REGISTER NODE** command to create an object client. Specify `TYPE=OBJECTCLIENT`. The script runs under the instance user ID.

What to do next

As a next step, you must register the IBM Spectrum Protect server as a repository server. For information about how to do this, see [“Registering a repository server as a backup storage provider”](#) on page 156. Once that is completed, you can create SLA policy jobs to copy data to the IBM Spectrum Protect server for standard storage or for archive to tape.

Registering a repository server as a backup storage provider

Add and register a repository server to enable IBM Spectrum Protect Plus to copy data to the server.

Before you begin

Configure the key and certificate that are required for the repository server. For instructions, see [“Adding an access key”](#) on page 166 and [“Adding a certificate”](#) on page 167.

For the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server.

Configure IBM Spectrum Protect Plus as an object client to the IBM Spectrum Protect server. The object client node transfers and stores copied data. After you complete the setup procedure, the wizard provides you with the endpoint for communicating with the object agent on the server, and the access ID, secret key, and certificate for connecting securely.

Important: Each IBM Spectrum Protect Plus server must be registered as its own object-client node in the IBM Spectrum Protect server.

Certificates can be obtained from the IBM Spectrum Protect server Operations Center by navigating to the following pane: **Server > Object Agent > Agent Certificate**. Alternatively, the certificate can be obtained from the IBM Spectrum Protect Plus appliance by running the following command: `openssl s_client -showcerts -connect <ip-address>:9000 </dev/null 2>/dev/null | openssl x509`

Copy retention settings are fully controlled through associated SLA policies in IBM Spectrum Protect Plus. IBM Spectrum Protect server copygroup retention settings are not used for copy operations.

Procedure

To add and register an IBM Spectrum Protect server as a backup storage provider, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Repository servers**.
2. Click **Add repository server** to open the **Add repository server** wizard.
3. Complete the fields on the **Repository server details** page, and then click **Next**:

Name

Enter a meaningful name to identify the repository server.

Hostname

Enter the high-level address (HLA) of the repository server object agent.

Tip: To retrieve the HLA, run the following command from the IBM Spectrum Protect server:

query server server_name format=detailed

where *server_name* specifies the object agent.

Port

Enter the communications port of the repository server.

Use existing key

Enable to select a previously entered key for the repository, and then select the key from the **Select a key** list.

If you do not select this option, complete the following fields to add a key:

Key name

Enter a meaningful name to identify the key.

Access key

Enter the access key.

Secret key

Enter the secret key.

Certificate

Select a method of associating a certificate with the resource. If copying the certificate, the BEGIN and END lines of text must be included.

Upload

Select and click **Browse** to locate the certificate, and then click **Upload**.

Copy and paste

Select to enter the name of the certificate, copy and paste the contents of the certificate, and then click **Create**.

Use existing

Select to use a previously uploaded certificate.

4. Review your selections, and then click **Submit**.

The repository server is added to the servers table.

What to do next

After you add a repository server, complete the following action:

Action	How to
Associate the repository server with the SLA policy that is used for the backup job.	To create an SLA policy, see “Creating an SLA policy for databases and file systems” on page 206. To modify an existing SLA policy, see “Editing an SLA policy” on page 210.

Related concepts

[“Configuration for copying or archiving data to IBM Spectrum Protect” on page 145](#)

If you are planning to copy or archive IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, there are three possible configurations. Choosing which one to configure depends on which scenario applies to your data protection needs. For each scenario, there are steps that are required in both the IBM Spectrum Protect Plus and IBM Spectrum Protect server environments to complete the setup.

Editing settings for a repository server

Edit the settings for a repository server provider to reflect changes in your cloud environment.

Procedure

To edit a repository server provider, complete the following steps:

1. In the In the navigation panel, click **System Configuration > Storage > Repository servers**.
2. Revise the settings for the repository server provider, and then click **Save**.
3. Select the server, and then click **Edit**.

The **Editing** pane is displayed.

Deleting a repository server

Delete a repository server provider to reflect changes in your environment. Ensure that the provider is not associated with any SLA policies before deleting the provider.

Before you begin



Attention: Deleting a repository server could result in a loss of data. Ensure that any data that you require is backed up before you delete the repository server.

Procedure

To delete a repository server provider, complete the following steps:

1. In the In the navigation panel, click **System Configuration > Storage > Repository servers**.
2. Select the server, and then click **Remove**.
3. Click **Yes** to delete the provider.

Managing sites

A *site* is an IBM Spectrum Protect Plus policy construct that is used to manage the placement of data in an environment.

A site can be physical, such as a data center, or logical, such as a department or organization. IBM Spectrum Protect Plus components are assigned to sites to localize and optimize data paths. An IBM Spectrum Protect Plus deployment always has at least one site per physical location.

By default, the IBM Spectrum Protect Plus environment has a Primary site and a Secondary site.

Adding a site

After you add a site to IBM Spectrum Protect Plus, you can assign backup storage servers to the site.

About this task

Options that are designated as "for vSnap server storage" apply only if a vSnap server is assigned to the site.

Procedure

To add a site, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Sites**.
2. Click **Add site** to open the **Add site** wizard.
3. Complete the fields on the **Site details** page, and then click **Next**:

Site name

Enter a site name.

Throughput throttle for vSnap server storage

To manage the network activity on a defined schedule, select **Enabled** and change the throughput for site replication and copy operations:

Throttle rate

Adjust the throughput rate:

- a. Change the numerical rate of throughput by clicking the up or down arrows.
- b. Select a unit for the throughput.

The default throughput is 100 MB per second.

The screenshot shows the 'Site details' configuration page. The 'Site name' field contains 'Secondary'. Under 'Throughput throttle', the 'Enabled' radio button is selected. The 'Throttle rate' is set to '525 MB per second'. Below this is a 'Throttle schedule' section with a grid for selecting times. The grid has rows for 'All', 'Sunday', 'Monday', 'Tuesday', 'Wednesday', 'Thursday', 'Friday', and 'Saturday', and columns for time slots from 0 to 34. Blue bars indicate enabled throttling: Sunday (8:00-8:59), Tuesday (7:00-7:59), Wednesday (8:00-8:59), Thursday (1:00-1:59, 8:00-8:59), Friday (8:00-8:59), and Saturday (4:00-4:59, 8:00-8:59). A legend on the right shows 'Enabled' as a blue bar and 'Disabled' as a white bar. At the bottom, there are 'Cancel' and 'Save' buttons.

Figure 17. Example of enabling different rates of throttling for different times to improve throughput

Throttle schedule for vSnap server storage

Select daily times for throttling, or select specific days and times for throttling. The time that is specified should be based on the local time of one or more vSnap servers that are assigned to the site.

Tip: To select a time, click a time slot in the table. The selected time slot is highlighted. To clear a time slot, click a highlighted time slot. To select the same time slot for every day of the week, click a time slot in the **All** row.

After you make your selections, throttling days and times are listed underneath the schedule table.

Resource allocation to storage

Do not change this option from the default unless instructed by IBM support.

Select one of the following options to determine how new resources, such as VMware, database, application, and container allocation backup data, is initially allocated to vSnap servers: by count, by free space, or a factor of both. For the **Primary** and **Secondary** site, the default value is **By count**. If you create a new site, the default value is **Balanced**.

Important: If you have set the **Resource allocation to storage** setting to **By free space** in IBM Spectrum Protect Plus 10.1.9, when you upgrade to IBM Spectrum Protect Plus 10.1.10, this is set

to **By count**. If you wish to use **By free space**, you will need to change to this setting after you upgrade.

By count

Data allocation is determined by number of resources on the storage. The vSnap server that contains the lowest number of resources is selected.

Balanced

Data allocation is determined equally by the resource count and free space on the storage.

By free space

Data allocation is determined by the amount of free space on the storage. The vSnap server that contains the highest amount of free space is selected.

Custom

Data allocation is determined by a custom factor by using a slider. The slider is set to allocate backup data as **By free space** as the default.

4. On the **Review** page, review your selections, and then click **Submit**.

Results

The site is displayed in the sites table and can be applied to new and existing backup storage servers.

Editing a site

You can revise site information to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit a site, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Sites**.
2. Select the site, and then click **Edit**.
The **Editing** pane is displayed.
3. Edit the following fields:

Site name

Enter a site name.

Throughput throttle for vSnap server storage

To manage the network activity on a defined schedule, select **Enabled** and change the throughput for site replication and copy operations:

Throttle rate

Adjust the throughput rate:

- a. Change the numerical rate of throughput by clicking the up or down arrows.
- b. Select a unit for the throughput.

The default throughput is 100 MB per second.

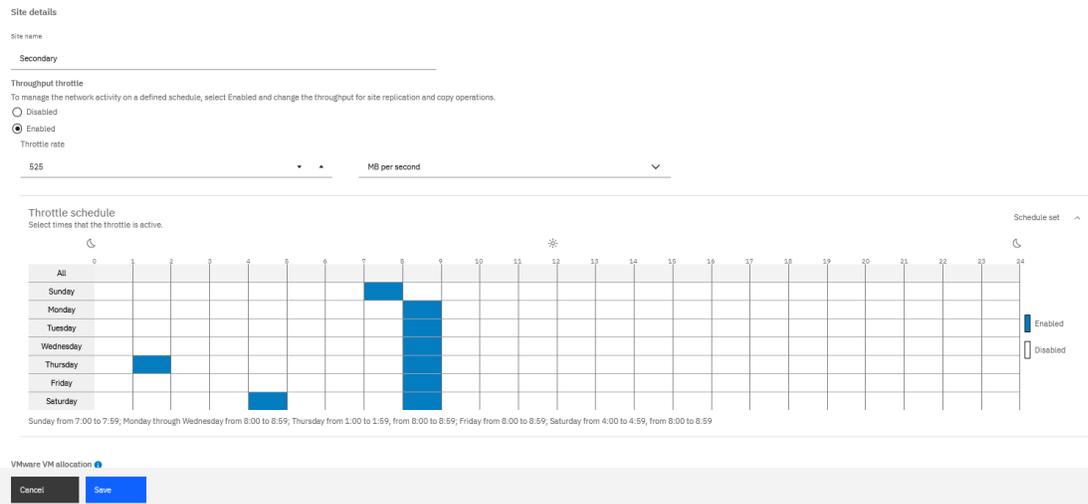


Figure 18. Example of enabling different rates of throttling for different times to improve throughput

Throttle schedule for vSnap server storage

Select daily times for throttling, or select specific days and times for throttling. The time that is specified should be based on the local time of one or more vSnap servers that are assigned to the site.

Tip: To select a time, click a time slot in the table. The selected time slot is highlighted. To clear a time slot, click a highlighted time slot. To select the same time slot for every day of the week, click a time slot in the **All** row.

After you make your selections, throttling days and times are listed underneath the schedule table.

Resource allocation to storage

Do not change this option from the default unless instructed by IBM support.

Select one of the following options to determine how new resources, such as VMware, database, application, and container allocation backup data, is initially allocated to vSnap servers: by count, by free space, or a factor of both. For the **Primary** and **Secondary** site, the default value is **By count**. If you create a new site, the default value is **Balanced**.

Important: If you have set the **Resource allocation to storage** setting to **By free space** in IBM Spectrum Protect Plus 10.1.9, when you upgrade to IBM Spectrum Protect Plus 10.1.10, this is set to **By count**. If you wish to use **By free space**, you will need to change to this setting after you upgrade.

By count

Data allocation is determined by number of resources on the storage. The vSnap server that contains the lowest number of resources is selected.

Balanced

Data allocation is determined equally by the resource count and free space on the storage.

By free space

Data allocation is determined by the amount of free space on the storage. The vSnap server that contains the highest amount of free space is selected.

Custom

Data allocation is determined by a custom factor by using a slider. The slider is set to allocate backup data as **By free space** as the default.

4. Click **Save** to commit the changes and close the pane.

Deleting a site

Delete a site when it becomes obsolete. Ensure that you reassign your backup storage to different sites before deleting the site.

Procedure

To delete a site, complete the following steps:

1. In the navigation panel, click **System Configuration > Storage > Sites**.
2. Select the site, and then click **Remove**.
3. Click **Yes** to delete the site.

Managing LDAP and SMTP servers

You can add a Lightweight Directory Access Protocol (LDAP) and Simple Mail Transfer Protocol (SMTP) server for use in the IBM Spectrum Protect Plus for use in user account and report features.

Related tasks

[“Creating a user account for an LDAP group” on page 466](#)

With IBM Spectrum Protect Plus, you can use a Lightweight Directory Access Protocol (LDAP) server to manage users. When you create an LDAP user account, you can add the user account to a user group.

[“Scheduling a report” on page 452](#)

You can schedule reports in IBM Spectrum Protect Plus to run at specific times.

Adding an LDAP server

You must add an LDAP server to create IBM Spectrum Protect Plus user accounts by using an LDAP group. These accounts allows users to access IBM Spectrum Protect Plus by using LDAP user names and passwords. Only one LDAP server can be associated with an instance of IBM Spectrum Protect Plus virtual appliance.

About this task

You can add a Microsoft Active Directory or OpenLDAP server. Note that OpenLDAP does not support the sAMAccountName user filter that is commonly used with Active Directory. Additionally, the **memberOf** option must be enabled on the OpenLDAP server.

Procedure

To register an LDAP server, complete the following steps:

1. In the navigation panel, click **System Configuration > LDAP/SMTP Server**.
2. In the **LDAP Servers** pane, click **Add LDAP Server**.
3. Populate the following fields in the **LDAP Servers** pane:

Host Address

The IP address of the host or logical name of the LDAP server.

Port

The port on which the LDAP server is listening. The typical default port is 389 for non TLS connections or 636 for TLS connections.

TLS

Enable the TLS option to establish a secure connection to the LDAP server.

Use existing user

Enable to select a previously entered user name and password for the LDAP server.

Bind Name

The bind distinguished name that is used for authenticating the connection to the LDAP server. IBM Spectrum Protect Plus supports simple bind.

Password

The password that is associated with the Bind Distinguished Name.

Base DN

The location where users and groups can be found.

User Filter

A filter to select only those users in the Base DN that match certain criteria. An example of a valid default user filter is `cn={0}`.

Tips:

- To enable authentication by using the **sAMAccountName** Windows user naming attribute, set the filter to `samaccountname={0}`. When this filter is set, users log in to IBM Spectrum Protect Plus by using only a user name. A domain is not included.
- To enable authentication using the user principal name (UPN) naming attribute, set the filter to `userprincipalname={0}`. When this filter is set, users log in to IBM Spectrum Protect Plus by using the `username@domain` format.
- To enable authentication by using an email address that is associated with LDAP, set the filter to `mail={0}`.

The **User Filter** setting also controls the type of user name that appears in the IBM Spectrum Protect Plus display of users.

User RDN

The relative distinguished path for the user. Specify the path where user records can be found. An example of a valid default RDN is `cn=Users`.

Group RDN

The relative distinguished path for the group. If the group is at a different level than the user path, specify the path where group records can be found.

4. Click **Save**.

Results

IBM Spectrum Protect Plus completes the following actions:

1. Confirms that a network connection is made.
2. Adds the LDAP server to the database.

After the SMTP server is added, the **Add LDAP Server** button is no longer available.

What to do next

If a message is returned indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

Related tasks

[“Creating a user account for an LDAP group” on page 466](#)

With IBM Spectrum Protect Plus, you can use a Lightweight Directory Access Protocol (LDAP) server to manage users. When you create an LDAP user account, you can add the user account to a user group.

Adding an SMTP server

You must add an SMTP server to send scheduled reports to email recipients. Only one SMTP server can be associated with a IBM Spectrum Protect Plus virtual appliance.

Before you begin

Beginning with IBM Spectrum Protect Plus 10.1.13, the user can select **Use TLS** button if their SMTP server supports TLS connections. When you enable the **Use TLS** button, the certification section is displayed on the **SMTP Servers** pane to import the SMTP certificate to the IBM Spectrum Protect Plus server.

Note: If the **Use TLS** button is not enabled, the SMTP server is registered in the traditional manner. This eliminates the need to use the certificate.

Procedure

To add an SMTP server, complete the following steps:

1. In the navigation panel, click **System Configuration > LDAP/SMTP Server**.
2. In the **SMTP Servers** pane, click **Add SMTP Server**.
3. Populate the following fields in the **SMTP Servers** pane:

Host Address

The IP address of the host, or the path and host name of the SMTP server.

Port

The communications port of the server that you are adding. The typical default port is 25 for non-TLS connections or 443 for TLS connections.

Username

The name that is used to access the SMTP server.

Password

The password that is associated with the user name.

Timeout

The email timeout value in milliseconds.

From Address

The address that is associated with email communications from IBM Spectrum Protect Plus.

Subject Prefix

The prefix to add to the email subject lines sent from IBM Spectrum Protect Plus.

Use TLS

Select this option when the SMTP server is using TLS.

Certificate

Tip: You must use the fully qualified domain name (FQDN) to validate the SMTP certificate. Using an IP address may fail to validate the SMTP certificate.

In the certificate field, select one of the following options to import the SMTP certificate to the IBM Spectrum Protect Plus server:

Use existing certificate

Click **Choose file** to select an existing certificate from the **Select a certificate** list.

Copy and paste

In the Enter certificate name field, enter a name for the certificate. Then, paste the contents of the certificate in the **Copy and paste certificate here** field and click **Create**.

Upload

- a. Download the file from the vCenter server to the local machine where you are running the browser.
- b. Click **Choose file** and search for the downloaded certificate in your system.
- c. Click **Upload**. This option is the default.

Get Certificate

Click **Get Certificate** to automatically retrieve the certificate that is associated with the SMTP Server that is specified in the hostname or IP address.

4. Click **Save**.

Results

IBM Spectrum Protect Plus completes the following actions:

1. Confirms that a network connection is made.
2. Adds the server to the database.

If a message is returned indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

To test the SMTP connection, click the **Test SMTP Server** button, then enter an e-mail address. Click **Send**. A test e-mail message is sent to the e-mail address to verify the connection.

After the SMTP server is added, the **Add SMTP Server** button is no longer available.

What to do next

Related tasks

[“Scheduling a report” on page 452](#)

You can schedule reports in IBM Spectrum Protect Plus to run at specific times.

Editing settings for an LDAP or SMTP server

Edit the settings for an LDAP or SMTP server to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit the settings for an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration > LDAP/SMTP Server**.
2. Click the edit icon  that is associated with the server.
The edit pane is displayed.
3. Revise the settings for the server, and then click **Save**.

Deleting an LDAP or SMTP server

Delete an LDAP or SMTP server when it becomes obsolete. Ensure that the server is not in use by IBM Spectrum Protect Plus before deleting the server.

Procedure

To delete an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration > LDAP/SMTP Server**.
2. Click the delete icon  that is associated with the server.
3. Click **Yes** to delete server.

Managing keys and certificates for connection to IBM Spectrum Protect Plus components

Keys and certificates are used in the IBM Spectrum Protect Plus environment to provide secure connections to IBM Spectrum Protect Plus components.

Keys, and in some environments certificates, are required to enable IBM Spectrum Protect Plus to connect to the following components:

Secondary backup storage

The cloud resources and repository servers that provide secondary backup storage require credentials to serve as copy destinations. Access keys and secret keys are provided by your cloud resource or repository server interface. These keys serve as the username and password of your copy destinations and allow them to be accessed by IBM Spectrum Protect Plus. Some copy destinations also require TLS certificates for additional data security. The TLS certificate can be a certificate that is issued by a certificate authority (CA).

Linux-based resources

You can add a Secure Shell (SSH) key to provide credentials for Linux-based resources on virtual machines managed by vCenter and Hyper-V, as well as Oracle, Db2, MongoDB, and SAP HANA application servers. SSH keys help to provide a secure connection between IBM Spectrum Protect Plus and target resources for file indexing and restore operations.

When you add a key or certificate to IBM Spectrum Protect Plus, the list of available keys and certificates is updated so that you can select a key or certificate as needed in the user interface.

For information about using a TLS certificate for secure connections to the IBM Spectrum Protect Plus user interface, see [“Uploading a TLS certificate” on page 171](#).

Adding an access key

Add an access key to provide cloud resource or repository server credentials.

Procedure

To add a key, complete the following steps:

1. Create your access key and secret key through the interface of the cloud resource or repository server. Make note of the access key and secret key.
2. In the navigation menu, click **System Configuration > Keys and Certificates**.
3. From the **Access Keys** section, click **Add Access Key**.
4. Complete the fields in the **Key Properties** pane:

Name

Enter a meaningful name to help identify the access key.

Access Key

Enter the access key of the cloud resource or repository server. For Microsoft Azure, enter the storage account name.

Secret Key

Enter the secret key of the cloud resource or repository server. For Microsoft Azure, enter the key from one of the key fields, either key1 or key2.

5. Click **Save**.

The key displays in the **Access Keys** table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing key** option.

Deleting an access key

Delete an access key when it becomes obsolete. Ensure that you reassign a new access key to your cloud resource or repository server.

Procedure

To delete an access key, complete the following steps:

1. In the navigation menu, click **System Configuration > Keys and Certificates**.
2. Click the delete icon  that is associated with an access key.
3. Click **Yes** to delete the access key.

Adding a certificate

Add a certificate to provide cloud resource or repository server credentials.

Procedure

To add a certificate, complete the following steps:

1. Export a certificate from your cloud resource or repository server.
2. In the navigation menu, click **System Configuration > Keys and Certificates**.
3. In the **Certificates** section, click **Add Certificate**.
4. Complete the fields in the **Certificate Properties** pane:

Type

Select the cloud resource or repository server type.

Certificate

Select a method to add the certificate:

Upload

Select to browse for the certificate locally.

Copy and paste

Select to enter the name of the certificate and copy and paste the contents of the certificate.

5. Click **Save**.

The key displays in the **Certificates** table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing certificate** option.

Deleting a certificate

Delete a certificate when it becomes obsolete. Ensure that you reassign a new certificate to your cloud resource or repository server.

Procedure

To delete a certificate, complete the following steps:

1. In the navigation menu, click **System Configuration > Keys and Certificates**.
2. Click the delete icon  that is associated with a certificate.
3. Click **Yes** to delete the certificate.

Adding an SSH key

You can add an SSH key to provide credentials for Linux-based and AIX-based resources on virtual machines managed by vCenter and Hyper-V, as well as Oracle, Db2, MongoDB, and SAP HANA application servers. SSH keys help to provide a secure connection between IBM Spectrum Protect Plus and target resources for file indexing and restore operations.

Before you begin

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The user account on the target resource that is used to generate the SSH key pair must have **sudo** privileges. This account, which will be assigned to IBM Spectrum Protect Plus, is known as the IBM Spectrum Protect Plus user agent (sppagent).
- If the environment includes virtual machines managed by vCenter, ensure that the latest VMware Tools are installed.

Procedure

To add a key, complete the following steps:

1. On the target resource, generate an SSH key by using the `ssh-keygen` command with the user account that will be assigned to IBM Spectrum Protect Plus. This account must have **sudo** privileges. For example, on an Oracle server, enter the following command in the terminal and follow the instructions:

```
ssh-keygen
```

If you use the default settings, two files are created in the specified directory: `id_rsa.pub` is the public key and `id_rsa` is the private key. The private key must be in PEM format. It may be necessary to explicitly use the `-m PEM` argument with `ssh-keygen` when generating the key pair.

2. When prompted enter the file name in which the key will be saved, enter a directory and file name. If you do not specify a directory and file name, the default is used:

```
/home/privileged_user/.ssh/id_rsa
```

where `privileged_user` is the account assigned to IBM Spectrum Protect Plus, `sppagent`. If a key with the default name already exists, this will be indicated with the message displayed below. Be careful not to overwrite preexisting keys if they are in use. Press **N** to enter a different file in which to save the key.

```
/home/<privileged user>/.ssh/id_rsa already exists.  
Overwrite (y/n)?
```

This procedure is based on the assumption that the key is saved in the default location using the default file name (`id_rsa`). If the key file is created using a different file name, use that file name in the steps that follow.

3. Supply a passphrase and press Enter. Otherwise, simply press Enter for no passphrase.
4. If a passphrase was supplied, enter it again. Press Enter.
5. Copy the contents of the `id_rsa.pub` key into the `authorized_keys` file. If the file already exists, append the public key to the `authorized_keys` file.

```
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
```

6. Assign the required privileges to the `authorized_keys` file by issuing the `chmod 600` command.

```
chmod 600 ~/.ssh/authorized_keys
```

7. Edit the `/etc/ssh/sshd_config` file to set the `PubkeyAuthentication` setting to `yes` by using a text editor. To ensure that the setting is not commented out, remove the number sign (`#`) if it appears at the beginning of the line.

```
sudo vi /etc/ssh/sshd_config
```

```
...  
PubkeyAuthentication yes  
...
```

8. Restart the SSH service on the target resource.

```
systemctl restart sshd
```

9. In the IBM Spectrum Protect Plus navigation panel, click **System Configuration > Keys and Certificates**.
10. From the **SSH Keys** section, click **Add SSH Key**.
11. Complete the fields in the **SSH Key Properties** pane:

Name

Enter a meaningful name to identify the SSH key.

User

Enter the user account that is associated with the target resource and SSH key. This is the user account used to generate the public and private keys in the previous steps.

Encrypted

Check this box if a passphrase was supplied when generating the public and private key.

Passphrase

This box is only displayed if the **Encrypted** check box is selected. If a passphrase was supplied when generating the public and private key, provide the passphrase in this box.

Private key

Copy and paste the private key into this box. This will be the key contained in the `id_rsa` file on the target resource. The file is similar to the following example:

```
cat ~/.ssh/id_rsa
```

```
-----BEGIN OPENSSH PRIVATE KEY-----  
ZRYtuinjaHx2mKgW4LnFqzlyAIIq5Amasi/J8/AAAFiFiP4GZYj+BmAAAAB3NzaC1yc2  
...  
...  
Q5ZqZ1Ec8N7dsAAAANDG9vckBVYnVudHVWQgECAwQFBg==  
-----END OPENSSH PRIVATE KEY-----
```

12. Click **Save**.

The key is displayed in the **SSH Keys** table and can be selected when you use a feature that requires credentials to access a resource with the **Key** option.

Deleting an SSH key

Delete an SSH key when it becomes obsolete. Ensure that you reassign a new SSH key to your resources.

Procedure

To delete an SSH key, complete the following steps:

1. In the navigation menu, click **System Configuration > Keys and Certificates**.
2. Click the delete icon  that is associated with an SSH key.

3. Click **Yes** to delete the access key.

Managing certificates for connection to the IBM Spectrum Protect Plus user interface

To establish secure connections to IBM Spectrum Protect Plus, you can upload the following TLS certificates depending on your environment and requirements: Hypertext Transfer Protocol Secure (HTTPS) and Lightweight Directory Access Protocol (LDAP).

An HTTPS certificate authority (CA) is required to establish a trusted connection to the IBM Spectrum Protect Plus user interface. You can start the IBM Spectrum Protect Plus user interface with the default self-signed certificate, but you will receive a browser notification that the certificate is not trusted.

An LDAP certificate is required if you are using LDAP authentication for IBM Spectrum Protect Plus users.

The following technotes provide introductory information for using certificates with IBM Spectrum Protect Plus:

HTTPS

[Technote 739663](#) provides information about using a HTTPS certificate that is issued by Microsoft Certificate Authority. However, you can use a certificate that is issued by another certificate authority (CA).

For HTTPS certificates, PEM encoded certificates with `.cer` or `.crt` extensions are supported.

LDAP

[Technote 791677](#) provides information about using an LDAP certificate.

For LDAP certificates, DER encoded certificates with `.cer` or `.crt` extensions are supported. If you are uploading an LDAP TLS certificate, ensure that IBM Spectrum Protect Plus has connectivity to the LDAP server and that the LDAP server is running.

CA certificate format example

The following example shows the format of a PEM encoded CA certificate. This example file contains a private key, a CA server certificate, one intermediate certificate, and a root certificate. The values in the private key and certificates are for example purposes.

```
# Private key
-----BEGIN PRIVATE KEY-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyz
-----END PRIVATE KEY-----

# CA server certificate
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
abcdefghijklmnopqrstuvwxyz
-----END CERTIFICATE-----

# Intermediate certificate
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyz
-----END CERTIFICATE-----

# Root certificate
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz
```

```
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZABCDEFGHIJKLMNOPQRSTUVWXYZ
-----END CERTIFICATE-----
```

Uploading a TLS certificate

You can upload TLS certificates to establish secure connections to IBM Spectrum Protect Plus.

Before you begin

To upload a TLS certificate, you must log on to IBM Spectrum Protect Plus as the superuser. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.

Procedure

To upload a TLS certificate, complete the following steps:

1. In the IBM Spectrum Protect Plus user interface, click the user menu  in the menu bar, and then click **Manage TLS certificates**.
2. Select the TLS certificate type: **HTTPS** or **LDAP**.
3. Click **Browse**, and select the certificate that you want to upload.
4. Click **Upload**.
5. Restart IBM Spectrum Protect Plus.

Testing network connectivity

The IBM Spectrum Protect Plus Service Tool tests host addresses and ports to determine if a connection can be established. You can use the Service Tool to verify whether a connection can be established between IBM Spectrum Protect Plus and a node.

You can run the Service Tool from the IBM Spectrum Protect Plus command line or remotely by using a .jar file. If a connection can be established, the tool returns a green check mark. If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

The tool provides guidance for the following error conditions:

- Timeout
- Connection refused
- Unknown host
- No route

Running the Service Tool from a command line

You can start the Service Tool from the IBM Spectrum Protect Plus virtual appliance command line interface and run the tool in a web browser. Then, you can use the Service Tool to verify network connectivity between IBM Spectrum Protect Plus and a node.

Procedure

1. Log in to the IBM Spectrum Protect Plus virtual appliance by using the serveradmin user ID and access the command line. Run the following command:

```
# sudo bash
```

2. Open port 9000 on the firewall by running the following command:

```
# firewall-cmd --add-port=9000/tcp
```

3. Run the tool by running the following command:

```
# java -Dserver.port=9000 -jar /opt/ECX/spp/public/assets/tool/ngxdd.jar
```

4. To connect to the tool, enter the following URL in a browser:

```
http://hostname:9000
```

where *hostname* specifies the IP address of the virtual machine where the application is deployed.

5. To specify the node to test, complete the following fields:

Host

The hostname or IP address of the node that you want to test.

Port

The connection port to test.

6. Click **Save**.

7. To run the tool, hover the cursor over the tool, and then click **Run**.

If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

8. Stop the tool by running the following command on the command line:

```
ctl-c
```

9. Protect your storage environment by resetting the firewall. Run the following commands:

```
# firewall-cmd --zone=public --remove-port=9000/tcp  
# firewall-cmd --runtime-to-permanent  
# firewall-cmd --reload
```

Note: If the `firewall-cmd` command is not available on your system, edit the firewall manually to add necessary ports and restart the firewall with `iptables`. For more information on editing firewall rules, see the **Firewall configuration with iptables** section here: [Examples of how to open firewall ports](#).

Running the Service Tool remotely

You can download the Service Tool as a .jar file from the IBM Spectrum Protect Plus user interface. Then, you can use the Service Tool to remotely test connectivity between IBM Spectrum Protect Plus and a node.

Procedure

1. In the IBM Spectrum Protect Plus user interface, click the user menu, and then click **Download Test Tool**.

A .jar file is downloaded to your workstation.

2. Launch the tool from a command-line interface. Java™ is only required on the system where the tool will be launched. Endpoints or target systems that are tested by the tool do not require Java.

The following command launches the tool in a Linux environment:

```
# java -jar -Dserver.port=9000 /<tool path >/ngxdd.jar
```

3. To connect to the tool, enter the following URL in a browser:

```
http://hostname:9000
```

where *hostname* specifies the IP address of the virtual machine where the application is deployed.

4. To specify the node to test, populate the following fields:

Host

The host name or IP address of the node that you want to test.

Port

The connection port to test.

5. Click **Save**.
6. To run the tool, hover the cursor over the tool, and then click the green **Run** button.
If a connection cannot be established, the error condition is displayed, along with possible causes and actions.
7. Stop the tool by issuing the following command on the command line:

```
ctl-c
```

Configuring global preferences

As the administrator, you can configure preferences that apply to all IBM Spectrum Protect Plus operations in the **Global Preferences** pane.

Before you begin

You must have administrator credentials to configure global preferences.

You can change the preference in the **Integration with other storage products** and **User interface** categories at any time.



Attention: Although you can modify the preference in the **Integrations with other storage products** and **User interface** categories, modify all other preferences only if absolutely necessary and only at the direction of IBM Support. Modifying global preferences can affect your storage environment. Preferences that require consultation with IBM Support are in the following categories: **Application, General, Job, Logging, Protection, and Security**.

About this task

Any changes that you make to parameter default values apply to all IBM Spectrum Protect Plus operations when you save the changes.

To return a global preference parameter to the default value, click the **x** icon in the setting field.

The figures in this task are for example purposes only and some values might differ from the default.

Procedure

To edit the values for any setting and apply them globally, complete the following steps:

1. In the navigation panel, click **System Configuration > Global Preferences**.
2. To enable access to IBM Spectrum Protect Operations Center from IBM Spectrum Protect Plus, edit the **IBM Spectrum Protect Operations Center URL** preference in the **Integration with other storage products** category.

The **IBM Spectrum Protect Operations Center URL** preference is the IP address of IBM Spectrum Protect Operations Center. The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment.

When this preference is set, the IBM Spectrum Protect icon  is active on the IBM Spectrum Protect Plus menu bar. When you initially set the URL for this preference or if you change it, you must log off and log back in for the preference to take effect in the user interface.

The URL is created during the Operations Center installation process. To obtain the Operations Center URL, contact the IBM Spectrum Protect system administrator.

Integration with other storage products

IBM Spectrum Protect Operations Center URL



https://myhost.mycompany.com:11090/oc

3. To set an automatic log out time for the IBM Spectrum Protect Plus user interface after a period of inactivity, edit the **Auto log out (minutes)** preference in the **User interface** category.

User interface

Auto logout (minutes)

30



4. To apply global application preferences, edit the settings in the **Application** category.

Application

Enable SQL Server and Oracle databases restored in test mode eligible for backup

Application inventory SSH timeout (hours)

12

Maximum volume size for backup target LUNs on Windows (TB)

256

Maximum backup retries (Kubernetes)

3

Maximum concurrent servers running backups

4



Overwrite existing SQL Server log backup from other SLA

Allow SQL database backup when transaction log backup chain is broken

Rename SQL files and folders when database is restored in production mode

You can edit the following application preferences:

Enable SQL Server and Oracle databases restored in test mode eligible for backup

Back up SQL Server databases and Oracle databases that were restored in test mode. When this option is selected, databases that were restored in test mode are available for selection in the SQL Backup pane, Oracle Backup pane, or ad hoc backup wizard.

Application inventory SSH timeout (hours)

The amount of time in hours after which the SSH session for an application inventory job will expire. The default setting is 12 hours with a maximum of 2048 hours.

Maximum volume size for backup target LUNs on Windows (TB)

The maximum size of the storage for a backup target.

Maximum backup retries (Kubernetes)

The maximum number of times that IBM Spectrum Protect Plus reattempts backup sessions for a copy backup job that contains multiple persistent volume claims (PVCs).

When multiple PVCs are involved in the same copy backup job, IBM Spectrum Protect Plus runs the backup operations as parallel jobs. To help prevent the backup sessions from timing out due to connection issues, specify the maximum number of times that IBM Spectrum Protect Plus reattempts the connections.

If the maximum number of retries is reached and connection failures still exist, only the PVC backups that were part of the failed sessions will be reported as failed.

Maximum concurrent servers running backups

The maximum number of concurrent application servers per backup session.

Overwrite existing SQL Server log backup from other SLA

Enable to permit overwriting of the existing SQL Server log backup that belongs to another SLA policy.

Allow SQL database backup when transaction log backup chain is broken

Run a database backup job when IBM Spectrum Protect Plus detects a break in the log backup chain for a database.

Rename SQL files and folders when database is restored in production mode

Rename associated SQL database data and log files during a restore job in production mode. This option applies only when a new database name is provided during an SQL database restore job.

5. To apply general preferences, edit the settings in the **General** category.

General	
Access log retention (days)	<input type="text" value="30"/>
Tools working folder on Linux guest	<input type="text" value="/tmp"/>
Tools working folder on Windows guest	<input type="text" value="c:\ProgramData"/>
Linux/AIX Clients Port (SSH) used for application and file indexing	<input type="text" value="22"/>
Windows Clients Port (WinRM) used for application and file indexing	<input type="text" value="5985"/>
Windows Clients (WinRM) deployment max retry on timeout	<input type="text" value="3"/>
IBM Spectrum Protect Plus Server NAT addresses	<input type="text"/>
IBM Spectrum Protect Plus Server IP Address 	<input type="text" value="9.11.67.160"/>

You can edit the following general preferences:

Access log retention (days)

Enter the number of days that the access log should be retained. The default setting is 30 days with a maximum setting of 1825 days (five years).

Tools working folder on Linux guest

The working folder for tools on Linux VM guests.

Tools working folder on Windows guest

The working folder for tools on Windows VM guests.

Linux/AIX Clients Port (SSH) used for application and file indexing

The SSH port that is used for application and file indexing on Linux and AIX clients.

Windows Clients Port (WinRM) used for application and file indexing

The Windows Remote Management (WinRM) port that is used for application and file indexing on Windows clients. Use port 5986 for encrypted communication. The default setting is 5985. You must sign out and then sign back in to IBM Spectrum Protect Plus to finalize this setting.

Note: When adding Windows client machines you must use port 5986 in order to make the **Get SSL certificate thumbprint** option visible.

Windows Clients (WinRM) deployment max retry on timeout

This is for Windows clients that utilize Windows Remote Management (WinRM). If a timeout occurs, the value determines the number of retries that should occur. The value can range from 0 to 5 with a default of 3.

IBM Spectrum Protect Plus Server NAT Addresses

The list of NAT IP addresses or FQDNs for the IBM Spectrum Protect Plus server.

IBM Spectrum Protect Plus Server IP Address

The list of available IP addresses for the IBM Spectrum Protect Plus server. The addresses are used for remote agent communication in NAT environments.

6. To apply job or logging preferences, edit the values in the **Job** or **Logging** categories.

The screenshot shows two sections of the configuration interface. The top section is titled "Job" and contains three input fields: "Job log retention (days)" with the value "60", "On-demand restore job retention (days)" with the value "3", and "Job notification status" with the value "failed". The bottom section is titled "Logging" and contains two options: "Enable logging IBM Spectrum Protect Plus alerts to the system log" with an unchecked checkbox, and "Enter which application types will skip downloading diagnostic logs (comma separated list, ex: sql,exch)" with an empty text input field.

You can edit the following job and logging preferences:

Job log retention (days)

The number of days to retain job logs before the logs are deleted.

On-demand restore job retention (days)

The number of days to retain on-demand restore jobs before the jobs, sessions, and logs are removed. If a value of 0 days is used, on-demand restore jobs are not removed. The default setting is 3 days.

Job notification status

The status level for sending alerts. Alerts are sent when a job is completed with the specified status. For example, if the job notification status is **failed**, when the failed status is reported for a job, an alert is sent.

Enable logging IBM Spectrum Protect Plus alerts to the system log

Include alerts that are generated by IBM Spectrum Protect Plus in the system log. After you enable this feature, you can search the system log to find alerts.

Enter which application types will skip downloading diagnostic logs (comma separated list, ex: sql, exch)

As part of each backup, IBM Spectrum Protect Plus collects a diagnostic log from the application server if the application agent generates a diagnostic log. You can add an application to this list to disable the downloading of the diagnostic logs for specific applications, even if they are generated by the application agent. Application types include `sql`, `oracle`, `db2`, `exch`, `mongo`, `saphana`, `kubernetesvol`, `office365`, `k8s`, `file`, and `openshift`. Multiple application types can be

entered and must be separated by commas with no intervening spaces. The default setting is to download all application agent diagnostic logs.

7. To apply protection preferences, edit the settings in the **Protection** category.

Protection	
Generate alert when time since last successful catalog backup is exceeded (hours)	24
File index overflow error (percentage)	90
File index overflow warning (percentage)	80
Number of seconds to wait before checking connection	1000
Number of times to check for valid connection	0
Temporary folder for file index zip files	/data2/filecatalog
Temporary folder for file indexing on Windows server	
Group VMs by	Count ▼
Number of VMs in group	5 ✕
Automatic removal of stale catalog entries	<input checked="" type="checkbox"/>
Enable backward compatibility cleanup (might increase the duration of the maintenance job)	<input type="checkbox"/>
Remove unselected VMs from SLA	<input type="checkbox"/>
Force the removal of replication relationship for last remaining snapshot	<input type="checkbox"/>
Select all replications for VMs and volumes	<input type="checkbox"/>
Target free space error (GB)	20
Target free space error (percentage)	1 ✕
Target free space warning (percentage)	5 ✕
Catalog object update count	50
Virtual machine backup status update interval (seconds)	300
Valid ESXi subnets for VADP proxy (comma separated, ex: 172.20.x.x)	
VADP proxy uses only HotAdd transport mode	<input type="checkbox"/>
vSnap auto disable deduplication when DDT size reaches resource limit	<input checked="" type="checkbox"/>
vSnap DDT size limit as percentage of total memory cache	80
vSnap DDT size limit in GB	50
Used space threshold on datastore or a volume before backup cannot take snapshots of a VM (percentage)	95
Backup wait timeout (seconds)	600
VMware communication timeout (seconds)	300

You can edit the following protection preferences:

Generate alert when time since last successful catalog backup is exceeded (hours)

The threshold in hours of when an alert should be generated if a successful catalog backup has not occurred. The alert is displayed on the IBM Spectrum Protect Plus product dashboard and in the **Alerts** window.

File index overflow error (percentage)

The threshold as a percentage for the file index count that when exceeded, results in an error in the job log. For example, if a value of 90 is specified, an error is displayed if the file index count reaches 90 percent of the allowed maximum of 2,147,483,647 entries.

File index overflow warning (percentage)

The threshold as a percentage for the file index count that when exceeded, results in a warning in the job log. For example, if a value of 80 is specified, a warning is displayed if the file index count reaches 80 percent of the allowed maximum of 2,147,483,647 entries.

Number of seconds to wait before checking connection

The amount of time that IBM Spectrum Protect Plus waits before checking the connection to a cloud object.

Number of times to check for valid connection

The number of times that IBM Spectrum Protect Plus checks for an available connection.

Temporary folder for file index zip files

The temporary folder for storing the compressed (.zip) files that contain the metadata for indexing. When the indexing is completed, the files are deleted.

Temporary folder for file indexing on Windows server

The temporary folder for storing the compressed (.zip) files that contain the metadata for indexing the Windows server. When the indexing is completed, the folder is deleted.

Group VMs by

Virtual machines can be grouped together. A group can be defined by a count of the VMs that are included in the group.

Number of VMs in group

For VM grouping, four VM groups are available and each VM group can have a maximum of five VMs. Each group corresponds to one destination volume (data stream). A maximum of 20 VMs (four data streams) can be grouped at a time based on size calculations.

Important: The default value for **Number of VMs in group** changed from 1 to 5 beginning in IBM Spectrum Protect Plus version 10.1.7.

If you are upgrading from version 10.1.6 and earlier and use the default value 1 for this option, the value is automatically changed to 5. If you have a custom value in this field, that value is maintained.

Automatic removal of stale catalog entries

Enables the checking and removal of stale catalog entries when compared with storage servers stored in MongoDB catalog. This action occurs during the execution of maintenance jobs. This will increase the overall time taken to complete the maintenance job.

Enable backward compatibility cleanup (might increase the duration of the maintenance job)

Enables the option to allow for the cleaning of both expired and unexpired snapshots that have reached their retention limit. Enabling this option might increase the duration of the maintenance job.

Remove unselected VMs from SLA

Enables the option to allow for VMware virtual machines (VMs) to be removed from the volume on vSnap server for corresponding SLA policies if the VM is no longer selected in the SLA policy. This will increase the overall time taken by maintenance and should only be enabled when required.

Force the removal of the replication relationship for last remaining snapshot

Remove an existing replication relationship for the last remaining snapshot that is set to expire and is locked.

Select all replications for VMs and volumes

Enables the capability for all replication jobs to select VMs and volumes from the beginning of the associated policy instead of selecting only from the last successful replication.

Target free space error (GB)

The size threshold in Gigabytes of remaining free space in the vSnap storage pool. Errors are displayed in the job log. For example, if a value of 100 is specified, an error is displayed if the vSnap storage pool has 100 GB or less of remaining free space. The error is generated when the threshold is reached based on the setting of either the size in this option or percentage as specified in **Target free space error (percentage)**. The resources will fail to back up to the affected vSnap storage pool if the threshold is reached. If this setting is changed, alerts are generated after you restart IBM Spectrum Protect Plus.

Target free space error (percentage)

The percentage threshold of remaining free space in the vSnap storage pool. Errors are displayed in the job log. For example, if a value of 5 is specified, an error is displayed if the vSnap storage pool has 5% or less of remaining free space. The error is generated when the threshold is reached based on the setting of either the percentage in this option or size as specified in **Target free space error (GB)**. The resources will fail to back up to the affected vSnap storage pool if the threshold is reached. If this setting is changed, alerts are generated after you restart IBM Spectrum Protect Plus.

Target free space warning (percentage)

The percentage threshold of remaining free space in the vSnap storage pool. Warnings are displayed in the job log. For example, if a value of 10 is specified, a warning is displayed if the vSnap storage pool has 10% or less of remaining free space. If this setting is changed, alerts are generated after you restart IBM Spectrum Protect Plus.

Catalog object update count

The count that you can set to limit how many objects are queried and updated in the catalog. For example, if the catalog includes 100 objects and the update count is 20, IBM Spectrum Protect Plus updates the catalog in five iterations.

Virtual machine backup status update interval (seconds)

The frequency at which messages about the progress of data transfer are updated in the job log.

Valid ESXi subnets for VADP proxy (comma separated, ex: 172.20.x.x)

Enter one or more valid subnets for VADP proxies. If more than one entry is provided, separate the subnets by commas with no intervening spaces.

VADP proxy uses only HotAdd transport mode

Use the HotAdd virtual disk transport method to connect the VMware IBM Spectrum Protect Plus virtual appliance with VADP proxies. If this option is enabled, VADP proxies will use HotAdd only without falling back to an alternate transport mode.

vSnap auto disable deduplication when DDT size reaches resource limit

The deduplication table (DDT) is enabled by default. When either of the threshold limits defined by disk space (gigabytes) or percentage is exceeded, vSnap data deduplication is disabled and an alert is displayed.

vSnap DDT size limit as percentage of total memory cache

The threshold as a percentage of the vSnap deduplication table (DDT) as compared to the total memory cache. The DDT is disabled when the vSnap auto disable option is selected and the defined threshold is exceeded.

vSnap DDT size limit in GB

The threshold, in gigabytes (GB), of the vSnap DDT. The DDT is disabled when the vSnap auto disable option is selected and the defined threshold is exceeded.

Used space threshold on datastore or a volume before backup cannot take snapshots of a VM (percentage)

The percentage of used space on a datastore or a volume that is the threshold before snapshots of a VM cannot be taken for backup.

Backup wait timeout (seconds)

The amount of time that IBM Spectrum Protect Plus waits for a backup job to finish before starting another backup job. If the backup job does not finish within the wait period, the job is timed out, and the next job begins.

VMware communication timeout (seconds)

The amount of time that IBM Spectrum Protect Plus waits for commands that are issued to connected vCenters to finish. If the operations do not finish within the specified amount of time, they are logged as errors. This setting applies only to VMware hypervisors.

- To apply a security preference, edit the setting in the **Security** category.



You can edit the following security preference:

Set minimum password length (characters)

The minimum length of passwords for IBM Spectrum Protect Plus. By default, the password has a minimum length of 8 characters, but you can specify a longer password. This value applies to all user accounts.

Configuring IBM Spectrum Protect Plus installed as a virtual appliance

There are some configuration tasks that apply only when IBM Spectrum Protect Plus is installed as a virtual appliance.

Logging on to the administrative console

Log on to the administrative console to review the configuration of the IBM Spectrum Protect Plus virtual appliance. Available information includes general system settings, network, and proxy settings.

Procedure

To log on to the administrative console, complete the following steps:

- From a supported browser, enter the following URL:

```
https://HOSTNAME:8090/
```

where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.

- In the login window, from the **Authentication Type** list, select one of the following authentication types:

Authentication Type	Logon information
IBM Spectrum Protect Plus	To log on as the IBM Spectrum Protect Plus superuser, enter the username and password. The IBM Spectrum Protect Plus super user is the user who is assigned the SUPERUSER role.
System	To log in as a system user, enter the serveradmin password.

Related concepts

[“System requirements ” on page 21](#)

Before you install IBM Spectrum Protect Plus, review the hardware and software requirements for the product and other components that you plan to install in the storage environment.

[“Managing roles” on page 461](#)

Roles define the actions that can be completed for the resources that are defined in a resource group. While a resource group defines the resources that are available to an account, a role sets the permissions to interact with the resources.

Logging on to the virtual appliance

Log on to the IBM Spectrum Protect Plus virtual appliance by using the vSphere Client to access the command line. You can access the command line in a VMware environment or in a Hyper-V environment.

Accessing the virtual appliance in VMware

In a VMware environment, log on to the IBM Spectrum Protect Plus virtual appliance through vSphere Client to access the command line.

Procedure

Complete the following steps to access the virtual appliance command line:

1. In vSphere Client, select the virtual machine where IBM Spectrum Protect Plus is deployed.
2. On the **Summary** tab, select **Open Console** and click in the console.
3. Select **Login**, and enter your user name and password. The default user name is `serveradmin` and the default password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).

What to do next

Enter commands to administer the virtual appliance. To log off, type `exit`.

Accessing the virtual appliance in Hyper-V

In a Hyper-V environment, log on to the IBM Spectrum Protect Plus virtual appliance through vSphere Client to access the command line.

Procedure

Complete the following steps to access the virtual appliance command line:

1. In Hyper-V Manager, select the virtual machine where IBM Spectrum Protect Plus is deployed.
2. Right-click the virtual machine and select **Connect**.
3. Select **Login**, and enter your user name and password. The default user name is `serveradmin` and the default password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).

What to do next

Enter commands to administer the virtual appliance. To log off, type `exit`.

Setting the time zone

Use the administrative console to set the time zone of the IBM Spectrum Protect Plus virtual appliance.

Procedure

To set the time zone, complete the following steps:

1. From a supported browser, enter the following URL:

```
https://HOSTNAME:8090/
```

where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.

2. In the login window, from the **Authentication Type** list, select one of the following authentication types:

Authentication Type	Login information
IBM Spectrum Protect Plus	To log in as the IBM Spectrum Protect Plus superuser, enter the username and password. The IBM Spectrum Protect Plus super user is the user who is assigned the SUPERUSER role.
System	To log in as a system user, enter the serveradmin password.

3. Click **System Management**.
4. In the **Timezone Management** section, select your time zone.
A message stating that the operation was successful displays. All IBM Spectrum Protect Plus logs and schedules will reflect the selected time zone. The selected time zone will also display on the IBM Spectrum Protect Plus virtual appliance when logged in with the user ID **serveradmin**.
5. Restart the IBM Spectrum Protect Plus virtual appliance from the administrative console.
6. When the IBM Spectrum Protect Plus virtual appliance has restarted, view the current time zone. Select **Product Information** from the main page of the administrative console and verify the updated time zone.

Adding virtual disks

You can add new virtual disks (hard disks) to your IBM Spectrum Protect Plus virtual appliance by using vCenter.

When you deploy the IBM Spectrum Protect Plus virtual appliance, you can deploy all virtual disks to one datastore that you specify at the time of deployment. You can add a disk within the virtual appliance and configure it as a Logical Volume Manager (LVM). You can then mount the new disk as a new volume or attach the new disk to the existing volumes within the virtual appliance.

Important: Do not add space or extend an existing volume for the IBM Spectrum Protect Plus virtual appliance.

You can review the disk partitions by using the **fdisk -l** command. You can review the physical volumes and the volume groups on the IBM Spectrum Protect Plus virtual appliance by using the **pvdisplay** and **vgdisplay** commands.

Adding a disk to the virtual appliance

Use the vCenter client to edit the settings of the virtual machine.

Before you begin

To run commands, you must connect to the command line for the IBM Spectrum Protect Plus virtual appliance by using Secure Shell (SSH) and log in with the user ID `serveradmin`. The default initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).

Procedure

To add a disk to an IBM Spectrum Protect Plus virtual appliance, complete the following steps from the vCenter client:

1. From the vCenter client, complete the following steps:
 - a) On the **Hardware** tab, click **Add**.
 - b) Select **Create a new virtual disk**.
 - c) Select the required disk size. In the **Location** section, select one of the following options:
 - To use the current datastore, select **Store with the virtual machine**.
 - To specify one or more datastores for the virtual disk, select **Specify a datastore or datastore cluster**. Click **Browse** to select the new datastores.
 - d) In the **Advanced Options** tab, leave the default values.
 - e) Review and save your changes.
 - f) Click the **Edit Settings** option for the virtual machine to view the new hard disk.
2. Add the new SCSI device without rebooting the virtual appliance. From the console of the IBM Spectrum Protect Plus appliance, issue the following commands:

```
sudo bash
```

Press Enter.

```
# for host in `ls /sys/class/scsi_host/`; do  
echo "- - -" > /sys/class/scsi_host/${host}/scan;  
done
```

Adding storage capacity from a new disk to the appliance volume

After you add a disk to the virtual appliance, you can attach the new disk to the existing volumes within the virtual appliance.

Before you begin

To run commands, you must connect to the console of the IBM Spectrum Protect Plus virtual appliance by using SSH and log in with the user ID `serveradmin`. The default initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 96](#).

About this task

You need to complete this task only if you want to add the storage capacity from a new disk to an existing appliance volume. If you added the disk as a new volume, you do not need to complete this task.

Procedure

To add storage capacity from a new disk to the appliance volume, complete the following steps from the console of the virtual appliance:

1. Complete the following steps to set up a partition for the new disk and set the partition to be of type Linux LVM:

- a) Open the new disk by using the **fdisk** command. For the command below, the disk `/dev/sdd` is used as an example. Use the **fdisk** command with the appropriate disk that is to be added.

```
[serveradmin@localhost ~]# fdisk /dev/sdd
```

The **fdisk** utility starts in interactive mode. Output similar to the following output is displayed:

```
Device contains neither a valid DOS partition table, nor Sun, SGI or
OSF disklabel
Building a new DOS disklabel with disk identifier 0xb1b293df.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.
Warning: invalid flag 0x0000 of partition table 4 will be corrected by
w(rite)
WARNING: DOS-compatible mode is deprecated. It's strongly recommended
to
switch off the mode (command 'c') and change display units to
sectors (command 'u').
Command (m for help):
```

- a) At the **fdisk** command line, enter the **n** subcommand to add a partition.

```
Command (m for help): n
```

The following command action choices are displayed:

```
Command (m for help): n
Command action
e extended
p primary partition (1-4)
```

- b) Enter the **p** command action to select the primary partition.
You are prompted for a partition number:

```
Command (m for help): n
Command action
e extended
p primary partition (1-4)
Partition number (1-4):
```

- c) At the partition number prompt, enter the partition number **1**.

```
Partition number (1-4): 1
```

The following prompt is displayed:

```
First cylinder (1-2610, default 1):
```

- d) Do not type anything at the First cylinder prompt. Press the **Enter** key.
The following output and prompt is displayed:

```
First cylinder (1-2610, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-2610, default 2610):
```

- e) Do not type anything in the Last cylinder prompt. Press the **Enter** key.
The following output is displayed:

```
Last cylinder, +cylinders or +size{K,M,G} (1-2610, default 2610):
Using default value 2610
Command (m for help):
```

- f) At the **fdisk** command line, enter the **t** subcommand to change a partition's system ID.

```
Command (m for help): t
```

You are prompted for a hex code that identifies the partition type:

```
Selected partition 1
Hex code (type L to list codes):
```

- g) At the Hex code prompt, enter the hex code 8e to specify the Linux LVM partition type.
The following output is displayed:

```
Hex code (type L to list codes): 8e
Changed system type of partition 1 to 8e (Linux LVM)
Command (m for help):
```

- h) At the **fdisk** command line, enter the **w** subcommand to write the partition table and to exit the **fdisk** utility.

```
Command (m for help): w
```

The following output is displayed:

```
Command (m for help): w (write table to disk and exit)
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

2. To review the changes to the disk, issue the **fdisk -l** command.
3. To review the current list of Physical Volumes (PV), issue the **pvdisplay** command.
4. To create a new Physical Volume (PV), issue the **pvcreate /dev/sdd1** command.
5. To view the new PV from /dev/sdd1, issue the **pvdisplay** command.
6. To review the Volume Group (VG), issue the **vgdisplay** command.
7. To add the Physical Volume (PV) to the Volume Group (VG) and increase the space of the VG, issue the following command:

```
vgextend data_vg /dev/sdd1
```

8. To verify that `data_vg` is extended, and that free space is available for logical volumes (or /data volume) to use, issue the **vgdisplay** command.
9. To review the Logical Volume (LV) /data volume, issue the **lvdisplay** command. The usage of the /data volume displays.
10. To add the space of the LV /data volume to the total volume capacity, issue the **lvextend** command.
In this example, 20 GB of space is being added to a 100 GB volume.

```
[serveradmin@localhost ~]# lvextend -l120gb -r /dev/data_vg/data
Size of logical volume data_vg/data changed from 100.00 GiB to 120.00 GiB .
Logical volume data successfully resized
resize2fs 1.41.12 (date)
Filesystem at /dev/mapper/data_vg-data is mounted on /data; on-line
resizing required
old desc_blocks = 7, new_desc_blocks = 8
Performing an on-line resize of /dev/mapper/data_vg-data to 31195136
(4k) blocks.
The filesystem on /dev/mapper/data_vg-data is now 31195136 blocks
long.
```

After you run the preceding command, the size of the /data volume is displayed in **lvdisplay** command output as 120 GB:

```
[serveradmin@localhost ~]# lvdisplay
--- Logical volume ---
LV Path: /dev/data_vg/data
LV Name: data
VG Name: data_vg
LV UUID: [uuid]
LV Write Access: read/write
LV Creation host, time localhost.localdomain, [date, time]
LV Status: available
# open: 1
LV Size: 120.00 GiB
Current LE: 30208
Segments : 2
Allocation inherit
Read ahead sectors: auto
- currently set to: 256
Block device: 253:1
[serveradmin@localhost ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/sda3 14G 2.6G 11G 20% /
tmpfs 16G 0 16G 0% /dev/shm
/dev/sda1 240M 40M 188M 18% /boot
/dev/mapper/data_vg-data
118G 6.4G 104G 6% /data
/dev/mapper/data2_vg-data2
246G 428M 234G 1% /data2
```

Resetting the serveradmin password

The `serveradmin` account is a system user account that is a preconfigured on IBM Spectrum Protect Plus and vSnap server. It is used to manage both of the virtual appliances. If you forget the `serveradmin` password, you must use the `root` account to reset the password. Because the password for the `root` account is not provided for deployments of IBM Spectrum Protect Plus or vSnap server, you must reset the `root` password and then reset the `serveradmin` through the virtual console.

About this task

Changing the `root` password will require that the IBM Spectrum Protect Plus virtual appliance be rebooted.

For IBM Spectrum Protect Plus, the preferred method is to avoid using the `root` account and instead use the `serveradmin` account for administration purposes.

Important: Pick a strong password that is a minimum of 15 characters in length and must contain at least one character from each of the classes (numbers, uppercase letters, lowercase letters, and other).

Procedure

1. Prepare to restart the IBM Spectrum Protect Plus virtual appliance by pausing all scheduled jobs. Then, wait for any running jobs to be completed.

Important: By ensuring that no jobs are running when you shut down the virtual appliance, you help to prevent possible issues.

2. Log in to the vSphere Client.
3. Restart the IBM Spectrum Protect Plus or vSnap server virtual appliance from the Actions menu by selecting Restart Guest OS.
4. Launch the web console. During the restart, the boot loader will appear.
5. Select the Red Hat Enterprise Linux (RHEL) version from the list and press the **e** key.
6. Locate the line that begins with `linux16 /vmlinuz`. In that line, locate `ro`.
7. Replace `ro` with the following string:

```
rw init=/sysroot/bin/sh
```

8. Press **Ctrl + X** to start in single user mode. A prompt appears.

9. Enter the chroot command:

```
:/# chroot /sysroot
```

10. Initiate the password change for the root account:

```
:/# passwd root
```

11. Enter the new password for the root account. You will be prompted to enter the password a second time.

12. Update the Security-Enhanced Linux parameters:

```
:/# touch /.autorelabel
```

13. Exit the current context with the exit command:

```
:/# exit
```

14. Restart the IBM Spectrum Protect Plus or vSnap server virtual appliance:

```
:/# reboot
```

15. Using secure shell (SSH) or the console, log in to the IBM Spectrum Protect Plus or vSnap server virtual appliance with the username `root` and the password that was created in a previous step.

16. Change the `serveradmin` password with the `passwd` command at the prompt:

```
:/# passwd serveradmin
```

17. Enter the new password for the `serveradmin` account. You will be prompted to enter the password a second time.

18. Close the SSH session to the IBM Spectrum Protect Plus or vSnap server virtual appliance.

19. Log in to IBM Spectrum Protect Plus or vSnap server using the `serveradmin` account with the newly created password to verify that the new password is set.

20. **For vSnap servers only:** Run the following command to re-sync the SMB/CIFS password with the new system password. When prompted to enter the password, specify the newly created password for the `serveradmin` account.

```
vsnap user resyncsmbpass --username serveradmin
```

Results

The `root` and `serveradmin` account passwords for the IBM Spectrum Protect Plus or vSnap server virtual appliance will be reset to the specified password.

Chapter 8. Managing SLA policies for backup operations

Service level agreement (SLA) policies, also known as backup policies, define parameters for backup jobs. These parameters include the frequency and retention period of backups, the backup location, and the option to replicate or copy backup data. You can use predefined SLA policies, or customize them to meet your needs.

The following default SLA policies are available. Each policy specifies a frequency and retention period for the backup. You can use these policies as they are or modify them. You can also create custom SLA policies.

Gold

This policy runs every 4 hours with a retention period of 1 week. For all supported resources except for Amazon EC2 instances and containers.

Silver

This policy runs daily with a retention period of 1 month. For all supported resources except for Amazon EC2 instances and container data.

Bronze

This policy runs daily with a retention period of 1 week. For all supported resources except for Amazon EC2 instances and container data.

EC2

To protect Amazon EC2 instances, this policy runs daily snapshot backups with a retention period of 31 days.

Container

To protect container data, this policy runs the following operations:

- Snapshot backups every 6 hours with a retention period of 1 day
- Copy backups daily with a retention period of 31 days.

To view and manage backup policies and to monitor the resources that are protected by policies, click **Manage Protection > Policy Overview** in the navigation panel.

If you edit an existing SLA policy by changing the standard object storage copy source, destination type, or target server options, the associated jobs will start a full base backup, not an incremental backup, during the next job run.

Selecting the primary and secondary backup location

During the process of creating an SLA policy, you must select a primary storage location. For some primary storage locations, you can also select the option to copy or archive data to a secondary storage location.

For a description of available primary and secondary backup storage options and instructions for adding storage systems in IBM Spectrum Protect Plus, see [“Managing backup storage” on page 136](#).

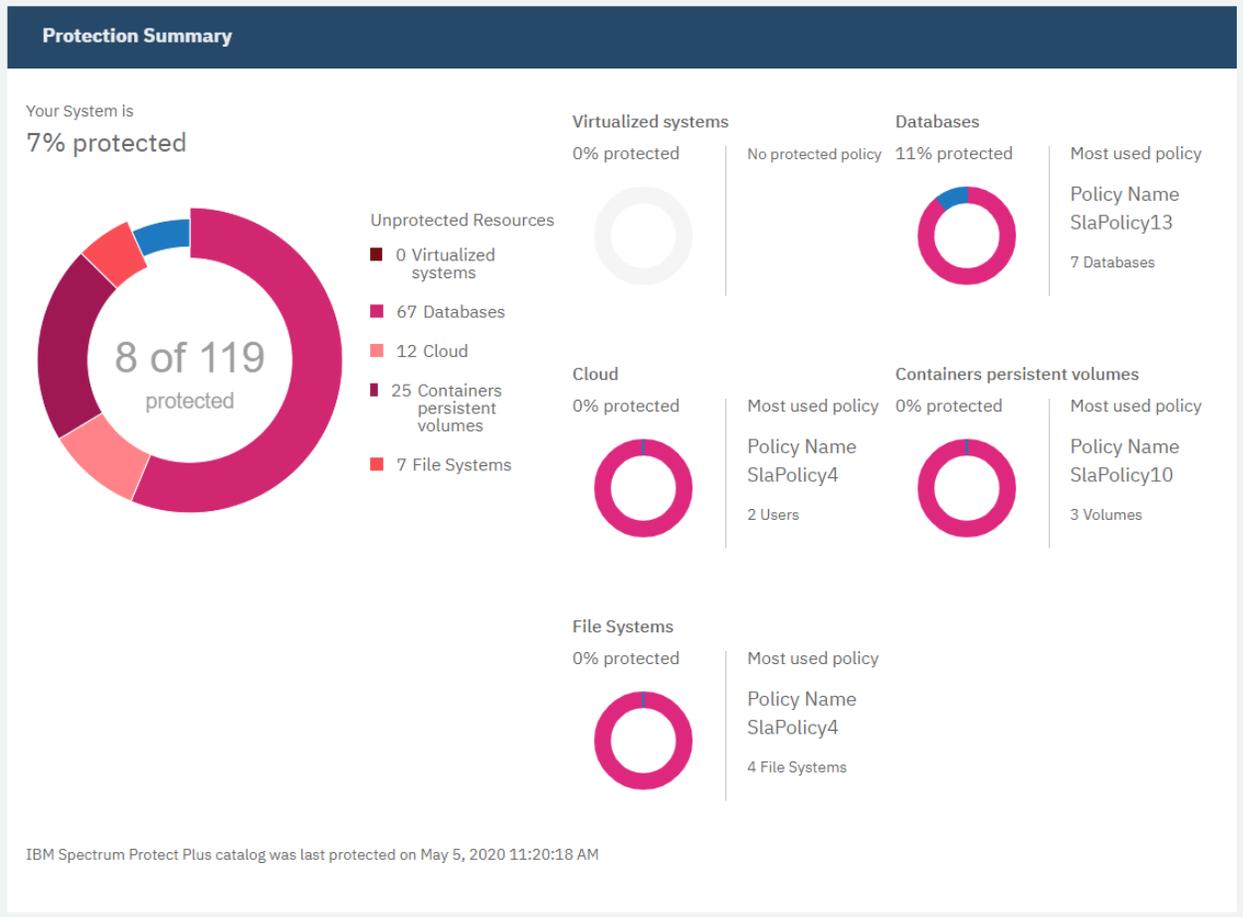
Protection Summary

You can view the protection status of the resources in your system in the **Protection Summary** pane.

The **Protection Summary** pane consists of donut charts that depict the number of protected resources versus the number of unprotected resources. For each type of resource, you can view the percentage of the resource that is protected, and the service level agreement (SLA) policy that is most frequently used for that resource.

To view the **Protection Summary** pane, from the navigation panel, click **Manage Protection > Policy Overview**.

Policy Overview



System

The **Your System** chart shows the total percentage of resources in your system that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of resources that are protected by IBM Spectrum Protect Plus. In the donut chart, the protected resources are represented by the blue line. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected resources.

Unprotected Resources

Shows the legend of unprotected resources. In the list, data is shown only for the types of resources that are managed by your instance of IBM Spectrum Protect Plus. If a type of resource is not managed by IBM Spectrum Protect Plus, the count is 0.

Virtualized systems

The **Virtualized systems** chart shows the percentage of virtualized system that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of virtualized systems that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected virtualized systems.

If no virtualized systems are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy

Shows the name of the most frequently used SLA policy, and the number of virtualized systems that are using that policy. If no virtualized systems are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy

This message is shown only when no virtualized systems are managed by IBM Spectrum Protect Plus.

Databases

The **Databases** chart shows the percentage of databases that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of databases that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected databases.

If no application databases are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy

Shows the name of the most frequently used SLA policy, and the number of databases that are using that policy. If no databases are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy

This message is shown only when no databases are managed by IBM Spectrum Protect Plus.

Cloud

The **Cloud** chart shows the percentage of cloud-based accounts, such as Microsoft Office 365 tenants, that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of cloud-based accounts that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected accounts.

If no cloud-based accounts are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy

Shows the name of the most frequently used SLA policy, and the number of accounts that are using that policy. If no cloud-based accounts are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy

This message is shown only when no cloud-based accounts are managed by IBM Spectrum Protect Plus.

Containers persistent volumes

Shows the percentage of persistent volumes that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of persistent volumes that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected persistent volumes.

If no persistent volumes are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy

Shows the name of the most frequently used SLA policy, and the number of persistent volumes that are using that policy. If no persistent volumes are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy

This message is shown only when no persistent volumes are managed by IBM Spectrum Protect Plus.

File Systems

Shows the percentage of file systems that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of file systems that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected file systems.

If no file systems are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy

Shows the name of the most frequently used SLA policy, and the number of file systems that are using that policy. If no file systems are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy

This message is shown only when no file systems are managed by IBM Spectrum Protect Plus.

Creating an SLA policy for the IBM Spectrum Protect Plus catalog

You can create custom service level agreement (SLA) policies to define options such as the backup location and schedule for the IBM Spectrum Protect Plus catalog. Depending on the back up location, you can also define options such as data retention, replication, and copy policies.

About this task

During the process of creating an SLA policy, you can select one of the following options for primary backup storage:

- A vSnap server
- Cloud storage

Ensure that the storage system that you want to associate with the SLA policy is configured in IBM Spectrum Protect Plus before you create the policy. For information about the adding a vSnap server or cloud storage, see [“Registering a vSnap server as a backup storage provider” on page 40](#) or [“Managing cloud storage” on page 138](#).

To optimize backup jobs for IBM Spectrum Protect Plus, create SLA policies specifically for backing up IBM Spectrum Protect Plus.

Creating an SLA policy for the IBM Spectrum Protect Plus catalog to back up to a vSnap server

You can create custom SLA policies that enable you to back up the IBM Spectrum Protect Plus catalog to a vSnap server.

Before you begin

Ensure that the storage system that you want to associate with the SLA policy is configured in **System Configuration > Storage > vSnap servers**. For information about the adding a vSnap server see [“Registering a vSnap server as a backup storage provider” on page 40](#).

About this task

If the catalog is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

If a snapshot replication task is started before an initial backup to a vSnap server is completed, errors in the job log indicate that no recovery points exist for the database. After the initial backup to the vSnap server is completed, run the replication task again to replicate the snapshots as configured in the SLA policy.

When copying data from a vSnap server to cloud storage, the most recent successfully completed snapshot will be copied.

Procedure

To create an SLA policy for the IBM Spectrum Protect Plus catalog, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.

2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **IBM Spectrum Protect Plus catalog** in the **Category** list.
4. Click **Catalog to vSnap** and then click **Next**.

The SLA policy options are displayed on the **Policy rules** page.

5. Complete the following options on the page and then click **Next**.
 - a) Enter a name that provides a meaningful description of the SLA policy.
 - b) Optional: Select the **Disable All Schedules** checkbox to disable all scheduling options in the policy.
 - c) In the **Backup Policy** section, set the following options for backup operations.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for backup operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for backing up data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage

Select this checkbox to back up data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

- d) In the **Log Backup Policy (SQL only)** section, set the following options for SQL log backup operations.

Enable Log Backup

Select this checkbox to enable the backing up of SQL transaction logs. These logs are used for recovery options such as point-in-time restore operations. If log backups are enabled for your backup jobs, transactions are continuously logged during the backup time. Notification is sent if any discontinuity is detected in log file backups.

Disable Schedule

Select this checkbox to create the log backup policy without defining a frequency or start time.

Repeats

Enter a frequency for log backup operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the log backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

- e) Under **Replication Policy**, set the following options to enable asynchronous replication from one vSnap server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in [“Configuring backup storage partners”](#) on page 47.

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this checkbox to create the replication policy without defining a frequency or start time.

Repeats

Enter a frequency for replication operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the replication operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for replicating data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage

Select this option to replicate data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

Same retention as source selection

Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Retention

Specify the retention period for the replicated data as a unit of time in days, months, or years.

- f) In the **Additional Copies** section, set the following options to copy data to standard object storage or archive object storage.

Standard object storage (incremental copy)

Select this option to copy data to cloud storage or to a repository server.

Data is backed up to the vSnap server for short term protection, and then copied to the selected cloud storage or repository server for longer-term protection. During the first copy of a backup volume, the snapshot is backed up in full. After the first copy of the base snapshot is completed, subsequent copies are incremental and capture cumulative changes since the last copy. Cloud or repository server restore operations can be performed from any available vSnap server.

Disable Schedule

Select this checkbox to create the copy policy without defining a frequency or start time.

Repeats

Enter a frequency for copy operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the copy operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Same retention as source selection

Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Restriction: This option and the **Retention** option are disabled if a server that uses write once read many (WORM) retention is selected in the **Target** field.

Retention

Specify the retention period for the copied snapshots as a unit of time in days, months, or years.

Source

Click the source for the copy operation:

Backup Policy Destination

The source for the copy operation is the target site that is defined in the **Backup Policy** section.

Replication Policy Destination

The source for the copy operation is the target site that is defined in the **Replication Policy** section.

This option is available only when **Backup Storage Replication** is selected.

Destination

Click **Cloud services** or **Repository servers**.

Target

Click the cloud storage system or repository server to which you want to copy data.

This list contains the secondary storage systems that you have added to IBM Spectrum Protect Plus. If you have not added secondary storage or want to add it, see [“Managing backup storage” on page 136](#) for information about the cloud storage systems and repository servers that are supported and how to add them to IBM Spectrum Protect Plus.

Archive object storage (full copy)

Select this option to archive data to cloud storage or to a repository server for long-term protection.

This operation provides a full image copy to the selected archival storage.

Disable Schedule

Select this checkbox to create the archive policy without defining a frequency or start time.

Repeats

Enter a frequency for archive operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the archive operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**.

You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Retention

Specify the retention period for the archive snapshots as a unit of time in days, months, or years.

Restriction: This option is disabled if a server that uses WORM retention is selected in the **Target** field.

Source

Click the source for the archive destination:

Backup Policy Destination

The source for the archive operation is the target site that is defined in the **Backup Policy** section.

Replication Policy Destination

The source for the archive operation is the target site that is defined in the **Replication Policy** section.

This option is available only when **Backup Storage Replication** is selected.

Destination

Click **Cloud services** or **Repository servers**.

Target

Click the cloud storage system or repository server to which you want to archive data.

Only cloud targets that have a defined archive bucket are shown in this list. To add an archive bucket for a cloud storage system, follow the instructions in [“Managing cloud storage” on page 138](#).

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, complete the following actions:

Action	How to
Assign user permissions to the SLA policy.	See “Creating a role” on page 463
Create a backup job definition that uses the SLA policy.	See “Backing up the IBM Spectrum Protect Plus catalog” on page 427 .

Related tasks

[“Editing an SLA policy” on page 210](#)

Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

[“Deleting an SLA policy” on page 211](#)

Delete an SLA policy when it becomes obsolete.

Creating an SLA policy for the IBM Spectrum Protect Plus catalog to back up to cloud storage

You can create custom SLA policies that enable you to back up the IBM Spectrum Protect Plus catalog to cloud storage.

Before you begin

Ensure that the storage system that you want to associate with the SLA policy is configured in **System Configuration > Storage > Cloud storage**. For information about the adding a vSnap server, see [“Managing cloud storage” on page 138](#).

About this task

If the catalog is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To create an SLA policy for the IBM Spectrum Protect Plus catalog, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **IBM Spectrum Protect Plus catalog** in the **Category** list.
4. Click **Catalog to object storage** and then click **Next**.

The SLA policy options are displayed on the **Policy rules** page.

5. Complete the following options on the page and then click **Next**.
 - a) Enter a name that provides a meaningful description of the SLA policy.
 - b) Set the following options for backup operations.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for backup operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you can select one or more days of the week. The **Start Time** applies to the selected days of the week.

Start Time

Enter the date and time that you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Destination

Click **Object Storage**.

Target Object Storage

Click the cloud storage system to which you want to archive data.

Only cloud targets that have a defined archive bucket are shown in this list. To add an archive bucket for a cloud storage system, follow the instructions in [“Managing cloud storage” on page 138](#).

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA policies** pane.

What to do next

After you create an SLA policy, complete the following actions:

Action	How to
Assign user permissions to the SLA policy.	See “Creating a role” on page 463
Create a backup job definition that uses the SLA policy.	See the backup topics in “Backing up the IBM Spectrum Protect Plus catalog” on page 427 .

Related tasks

[“Editing an SLA policy” on page 210](#)

Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

[“Deleting an SLA policy” on page 211](#)

Delete an SLA policy when it becomes obsolete.

Creating an SLA policy for hypervisors

You can create custom service level agreement (SLA) policies to define options such as the backup location and schedule for hypervisor virtual machines (VMs). Depending on the back up location, you can also define options such as data retention, replication, and copy policies.

Before you begin

Ensure that the storage system that you want to associate with the SLA policy is configured in IBM Spectrum Protect Plus before you create the policy. For information about the adding a vSnap server or OSSM server, see [“Registering a vSnap server as a backup storage provider” on page 40](#) or [“Adding the Open Snap Store Manager server as a backup storage provider” on page 120](#).

About this task

During the process of creating an SLA policy, you can select one of the following options for primary backup storage:

- A vSnap server (VMware or Microsoft Hyper-V VMs)
- An Open Snap Store Manager (OSSM) server (only VMware VMs)

To create an SLA policy to back up to a vSnap server, see the instructions in [“Creating an SLA policy for hypervisor backup to a vSnap server” on page 198](#)

To create an SLA policy to back up to an OSSM server, see the instructions in [“Creating an SLA policy for VMware backup to the OSSM storage server” on page 127](#)

Creating an SLA policy for hypervisor backup to a vSnap server

You can create custom SLA policies that enable you to back up hypervisor in your environment to a vSnap server.

Before you begin

Ensure that the storage system that you want to associate with the SLA policy is configured in **System Configuration > Storage > vSnap servers**. For information about the adding a vSnap server see [“Registering a vSnap server as a backup storage provider” on page 40](#).

About this task

If a virtual machine is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

If a snapshot replication task is started before an initial backup to a vSnap server is completed, errors in the job log indicate that no recovery points exist for the database. After the initial backup to the vSnap server is completed, run the replication task again to replicate the snapshots as configured in the SLA policy.

When copying data from a vSnap server to cloud storage, the most recent successfully completed snapshot will be copied.

Procedure

To create an SLA policy for hypervisor that back up to a vSnap server, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **Virtualized systems** in the **Category** list.
4. Click **Tiered vSnap** and then click **Next**.

The SLA policy options are displayed on the **Policy rules** page.

5. Complete the following options on the page and then click **Next**.
 - a) Enter a name that provides a meaningful description of the SLA policy.
 - b) Optional: Select the **Disable All Schedules** checkbox to disable all scheduling options in the policy.
 - c) In the **Backup Policy** section, set the following options for backup operations.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for backup operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for backing up data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage

Select this checkbox to back up data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

- d) In the **Log Backup Policy (SQL only)** section, set the following options for SQL log backup operations.

Enable Log Backup

Select this checkbox to enable the backing up of SQL transaction logs. These logs are used for recovery options such as point-in-time restore operations. If log backups are enabled for your backup jobs, transactions are continuously logged during the backup time. Notification is sent if any discontinuity is detected in log file backups.

Disable Schedule

Select this checkbox to create the log backup policy without defining a frequency or start time.

Repeats

Enter a frequency for log backup operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the log backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

- e) Under **Replication Policy**, set the following options to enable asynchronous replication from one vSnap server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in [“Configuring backup storage partners”](#) on page 47.

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this checkbox to create the replication policy without defining a frequency or start time.

Repeats

Enter a frequency for replication operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the replication operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for replicating data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage

Select this option to replicate data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

Same retention as source selection

Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Retention

Specify the retention period for the replicated data as a unit of time in days, months, or years.

- f) In the **Additional Copies** section, set the following options to copy data to standard object storage or archive object storage.

Standard object storage (incremental copy)

Select this option to copy data to cloud storage or to a repository server.

Data is backed up to the vSnap server for short term protection, and then copied to the selected cloud storage or repository server for longer-term protection. During the first copy of a backup volume, the snapshot is backed up in full. After the first copy of the base snapshot is completed, subsequent copies are incremental and capture cumulative changes since the last copy. Cloud or repository server restore operations can be performed from any available vSnap server.

Disable Schedule

Select this checkbox to create the copy policy without defining a frequency or start time.

Repeats

Enter a frequency for copy operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the copy operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Same retention as source selection

Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Restriction: This option and the **Retention** option are disabled if a server that uses write once read many (WORM) retention is selected in the **Target** field.

Retention

Specify the retention period for the copied snapshots as a unit of time in days, months, or years.

Source

Click the source for the copy operation:

Backup Policy Destination

The source for the copy operation is the target site that is defined in the **Backup Policy** section.

Replication Policy Destination

The source for the copy operation is the target site that is defined in the **Replication Policy** section.

This option is available only when **Backup Storage Replication** is selected.

Destination

Click **Cloud services** or **Repository servers**.

Target

Click the cloud storage system or repository server to which you want to copy data.

This list contains the secondary storage systems that you have added to IBM Spectrum Protect Plus. If you have not added secondary storage or want to add it, see [“Managing backup storage” on page 136](#) for information about the cloud storage systems and repository servers that are supported and how to add them to IBM Spectrum Protect Plus.

Archive object storage (full copy)

Select this option to archive data to cloud storage or to a repository server for long-term protection.

This operation provides a full image copy to the selected archival storage.

Disable Schedule

Select this checkbox to create the archive policy without defining a frequency or start time.

Repeats

Enter a frequency for archive operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the archive operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Retention

Specify the retention period for the archive snapshots as a unit of time in days, months, or years.

Restriction: This option is disabled if a server that uses WORM retention is selected in the **Target** field.

Source

Click the source for the archive destination:

Backup Policy Destination

The source for the archive operation is the target site that is defined in the **Backup Policy** section.

Replication Policy Destination

The source for the archive operation is the target site that is defined in the **Replication Policy** section.

This option is available only when **Backup Storage Replication** is selected.

Destination

Click **Cloud services** or **Repository servers**.

Target

Click the cloud storage system or repository server to which you want to archive data.

Only cloud targets that have a defined archive bucket are shown in this list. To add an archive bucket for a cloud storage system, follow the instructions in [“Managing cloud storage”](#) on page 138.

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, complete the following actions:

Action	How to
Assign user permissions to the SLA policy.	See “Creating a role” on page 463
Create a backup job definition that uses the SLA policy.	See the backup topics in Chapter 9, “Protecting virtualized systems,” on page 213.

Related concepts

[“Replicate backup-storage data ”](#) on page 13

When you enable replication of backup data, data from one vSnap server is asynchronously replicated to another vSnap server. For example, you can replicate backup data from a vSnap server on a primary site to a vSnap server on a secondary site.

[“Copying snapshots to secondary backup storage”](#) on page 14

If your primary backup storage is a vSnap server, you can copy snapshots from the primary backup storage to secondary storage for longer-term data protection. Secondary storage is not available for container data that is backed up to cloud storage.

Related tasks

[“Creating an SLA policy for Amazon EC2 instances” on page 205](#)

You can create custom service level agreement (SLA) policies to define snapshot retention and frequency policies that are specific to Amazon EC2 instances.

Creating an SLA policy for VMware backup to the OSSM storage server

You can create the custom SLA policies that enable you to back up VMware data to the Open Snap Store Manager (OSSM) storage server.

Before you begin

Ensure that the storage system that you want to associate with an SLA policy is configured in **System Configuration > Storage > OSSM**. For information about adding the OSSM storage server, see [“Adding the Open Snap Store Manager server as a backup storage provider” on page 120](#).

About this task

If a virtual machine is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To create an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **Virtualized systems** in the **Category** list.
4. Click **OSSM** and then click **Next**.

The SLA policy options are displayed on the **Policy rules** page.

5. Complete the following options on the page, and then click **Next**.
 - a) In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
 - b) Optional: Select **Disable all Schedules** checkbox to disable all scheduling options in the policy.
 - c) In the **Backup Policy** section, set the following options for backup operations.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for backup operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you might select one or more days of the week. The **Start Time** applies to the selected days of the week.

Start Time

Enter the date and time when you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for backing up the data.

Only sites that are associated with an OSSM storage server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with an OSSM storage server, are not shown.

- d) In the **Replication Policy** section, set the following options to enable asynchronous replication from one OSSM server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in [“Configuring backup storage partners for Open Snap Store Manager”](#) on page 129.

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this checkbox to create the replication relationship without defining a frequency or start time.

Repeats

Enter a frequency for replication operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for replicating data.

A site can only contain one OSSM storage server.

Only sites that are associated with a OSSM server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a OSSM server, are not shown.

Same retention as source selection

Select this option to use the same retention policy as the source OSSM server. To set a different retention policy, clear this option and set a different policy.

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, refer to [“Managing Jobs for VMware backups”](#) on page 129 to run a backup and restore jobs.

Related information

[Editing an SLA policy](#)

[Deleting an SLA policy](#)

Creating an SLA policy for Amazon EC2 instances

You can create custom service level agreement (SLA) policies to define snapshot retention and frequency policies that are specific to Amazon EC2 instances.

About this task

When a scheduled backup job runs, a snapshot of the instance is created at the frequency that is defined by the snapshot policy.

If an instance is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To create an SLA policy for your instances, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **Virtualized systems** in the **Category** list.
4. Click **Cloud snapshot** and then click **Next**.

The SLA policy options are displayed on the **Policy rules** page.

5. Complete the following options on the page and then click **Next**.

- a) Enter a name that provides a meaningful description of the SLA policy.
- b) In the **Snapshot Protection** section, set the following options for snapshot operations and then click **Next**:

Retention

Specify the retention period for the snapshots.

Disable Schedule

Select this checkbox to create the snapshot policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for snapshot operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you can select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time when you want the snapshot operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Snapshot Prefix

Enter a prefix to add to the beginning of snapshot names. Prefixes can help you organize and easily identify snapshots.

For example, if you entered the prefix "daily_", all snapshot names that are created with this SLA policy will begin with "daily_".

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, complete the following actions:

- Assign user permissions to the SLA policy. For instructions, see [“Creating a role” on page 463](#).
- Create a backup job definition that uses the SLA policy. For instructions, see [“Backing up Amazon EC2 data” on page 259](#).

Related tasks

[“Editing an SLA policy” on page 210](#)

Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

[“Deleting an SLA policy” on page 211](#)

Delete an SLA policy when it becomes obsolete.

Creating an SLA policy for databases and file systems

You can create custom service level agreement (SLA) policies to define backup frequency, retention, replication, and copy policies that are specific for databases and file systems.

Before you begin

Ensure that the storage system that you want to associate with the SLA policy is configured in **System Configuration > Storage > vSnap servers**. For information about the adding a vSnap server see [“Registering a vSnap server as a backup storage provider” on page 40](#).

About this task

If a virtual machine is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

If a snapshot replication task is started before an initial backup to a vSnap server is completed, errors in the job log indicate that no recovery points exist for the database. After the initial backup to the vSnap server is completed, run the replication task again to replicate the snapshots as configured in the SLA policy.

When copying data from a vSnap server to cloud storage, the most recent successfully completed snapshot will be copied.

Procedure

To create an SLA policy for databases and file systems, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy** to open the **Add SLA Policy** wizard.
3. Select **Applications and databases** in the **Category** list.
4. Click **Tiered vSnap** and then click **Next**.
The SLA policy options are displayed on the **Policy rules** page.
5. Complete the following options on the page and then click **Next**.
 - a) Enter a name that provides a meaningful description of the SLA policy.
 - b) Optional: Select the **Disable All Schedules** checkbox to disable all scheduling options in the policy.
 - c) In the **Backup Policy** section, set the following options for backup operations.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

Repeats

Enter a frequency for backup operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for backing up data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage

Select this checkbox to back up data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

- d) In the **Log Backup Policy (SQL only)** section, set the following options for SQL log backup operations.

Enable Log Backup

Select this checkbox to enable the backing up of SQL transaction logs. These logs are used for recovery options such as point-in-time restore operations. If log backups are enabled for your backup jobs, transactions are continuously logged during the backup time. Notification is sent if any discontinuity is detected in log file backups.

Disable Schedule

Select this checkbox to create the log backup policy without defining a frequency or start time.

Repeats

Enter a frequency for log backup operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the log backup operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

- e) Under **Replication Policy**, set the following options to enable asynchronous replication from one vSnap server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in [“Configuring backup storage partners” on page 47](#).

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this checkbox to create the replication policy without defining a frequency or start time.

Repeats

Enter a frequency for replication operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the replication operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site

Select the target backup site for replicating data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage

Select this option to replicate data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

Same retention as source selection

Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Retention

Specify the retention period for the replicated data as a unit of time in days, months, or years.

- f) In the **Additional Copies** section, set the following options to copy data to standard object storage or archive object storage.

Standard object storage (incremental copy)

Select this option to copy data to cloud storage or to a repository server.

Data is backed up to the vSnap server for short term protection, and then copied to the selected cloud storage or repository server for longer-term protection. During the first copy of a backup volume, the snapshot is backed up in full. After the first copy of the base snapshot is completed, subsequent copies are incremental and capture cumulative changes since the last copy. Cloud or repository server restore operations can be performed from any available vSnap server.

Disable Schedule

Select this checkbox to create the copy policy without defining a frequency or start time.

Repeats

Enter a frequency for copy operations. Choose from **Subhourly, Hourly, Daily, Weekly, Monthly, or Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the copy operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Same retention as source selection

Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Restriction: This option and the **Retention** option are disabled if a server that uses write once read many (WORM) retention is selected in the **Target** field.

Retention

Specify the retention period for the copied snapshots as a unit of time in days, months, or years.

Source

Click the source for the copy operation:

Backup Policy Destination

The source for the copy operation is the target site that is defined in the **Backup Policy** section.

Replication Policy Destination

The source for the copy operation is the target site that is defined in the **Replication Policy** section.

This option is available only when **Backup Storage Replication** is selected.

Destination

Click **Cloud services** or **Repository servers**.

Target

Click the cloud storage system or repository server to which you want to copy data.

This list contains the secondary storage systems that you have added to IBM Spectrum Protect Plus. If you have not added secondary storage or want to add it, see [“Managing backup storage” on page 136](#) for information about the cloud storage systems and repository servers that are supported and how to add them to IBM Spectrum Protect Plus.

Archive object storage (full copy)

Select this option to archive data to cloud storage or to a repository server for long-term protection.

This operation provides a full image copy to the selected archival storage.

Disable Schedule

Select this checkbox to create the archive policy without defining a frequency or start time.

Repeats

Enter a frequency for archive operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

Start Time

Enter the date and time that you want the archive operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Retention

Specify the retention period for the archive snapshots as a unit of time in days, months, or years.

Restriction: This option is disabled if a server that uses WORM retention is selected in the **Target** field.

Source

Click the source for the archive destination:

Backup Policy Destination

The source for the archive operation is the target site that is defined in the **Backup Policy** section.

Replication Policy Destination

The source for the archive operation is the target site that is defined in the **Replication Policy** section.

This option is available only when **Backup Storage Replication** is selected.

Destination

Click **Cloud services** or **Repository servers**.

Target

Click the cloud storage system or repository server to which you want to archive data.

Only cloud targets that have a defined archive bucket are shown in this list. To add an archive bucket for a cloud storage system, follow the instructions in [“Managing cloud storage”](#) on page 138.

6. Review your selections, and then click **Submit**.

The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, complete the following actions:

Action	How to
Assign user permissions to the SLA policy.	See “Creating a role” on page 463
Create a backup job definition that uses the SLA policy.	See the backup topics in Chapter 12, “Protecting databases,” on page 289 and Chapter 10, “Protecting file systems,” on page 267.

Related concepts

[“Replicate backup-storage data ”](#) on page 13

When you enable replication of backup data, data from one vSnap server is asynchronously replicated to another vSnap server. For example, you can replicate backup data from a vSnap server on a primary site to a vSnap server on a secondary site.

[“Copying snapshots to secondary backup storage”](#) on page 14

If your primary backup storage is a vSnap server, you can copy snapshots from the primary backup storage to secondary storage for longer-term data protection. Secondary storage is not available for container data that is backed up to cloud storage.

Related tasks

[“Creating an SLA policy for Amazon EC2 instances”](#) on page 205

You can create custom service level agreement (SLA) policies to define snapshot retention and frequency policies that are specific to Amazon EC2 instances.

Editing an SLA policy

Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click the edit icon  that is associated with a policy.
The **Edit SLA policy name** pane is displayed.
3. Edit the policy options, and then click **Save**.

Deleting an SLA policy

Delete an SLA policy when it becomes obsolete.

Before you begin

Ensure that there are no jobs that are associated with the SLA policy.

When the last remaining resource is disassociated from an SLA policy, IBM Spectrum Protect Plus detects that no resources are associated with the SLA policy. Any jobs that use the SLA policy will be changed to the **Held** state automatically. You are not able to delete the SLA policy because of dependent jobs. To delete an SLA policy, navigate to **Jobs and Operations** and click the **Schedule** tab. Verify that the jobs are in the **Held** state and then delete the jobs. Then you can delete the SLA policy.

If a new resource is added to the SLA policy and dependent jobs are not deleted, the job ID for each job will be maintained, but you will have to release any jobs in the **Held** state for them to continue.

Procedure

To delete an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection > Policy Overview**.
2. Click the delete icon  that is associated with an SLA policy.
3. Click **Yes** to delete the policy.

Chapter 9. Protecting virtualized systems

You must register the virtualized systems that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the resources that are associated with the systems.

Virtualized systems refers to VMware and Microsoft Hyper-V hypervisors and Amazon EC2 instances.

Backing up and restoring VMware data

To protect VMware data, first add vCenter Server instances in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for the content of the instances.

Ensure that your VMware environment meets the system requirements in [“Hypervisor \(Microsoft Hyper-V and VMware\) and cloud instance \(Amazon EC2\) backup and restore requirements ”](#) on page 21.

Support for VMware tags

IBM Spectrum Protect Plus supports VMware virtual machine tags. Tags are applied in vSphere and allow users to assign metadata to virtual machines. When applied in vSphere and added to the IBM Spectrum Protect Plus inventory, virtual machine tags can be viewed through the **View > Tags & Categories** filter when you create a job definition. For more information about VMware tagging, see [Tagging Objects](#). You must assign tags at the VM guest level for them to be utilized for backup exclusion rules based on tags or to be used as filtering for reports in IBM Spectrum Protect Plus.

Support for encryption

Backing up and restoring encrypted virtual machines is supported in vSphere 6.5 environments and later. Encrypted virtual machines can be backed up and restored at the virtual-machine level to their original location. If you are restoring a virtual machine to an alternative location, the encrypted virtual machine is restored without encryption, and must be encrypted manually by using the vCenter Server after the restore operation is completed.

The following vCenter Server privileges are required to enable operations for encrypted virtual machines:

- Cryptographer.Access
- Cryptographer.AddDisk
- Cryptographer.Clone

Note: An NFS volume may be mounted to any number of datacenters that belong to the same vCenter. If an NFS volume is mounted on more than one datacenter, vCenter treats the same volume as two different datastores. IBM Spectrum Protect Plus treats this as a single datastore and combines all of the VMs and VMDKs residing on the datastore from all of the datacenters on which the datastore is mounted. Any SLA selection against this datastore will cause all of the VMs from the different datacenters to be backed up or restored in IBM Spectrum Protect Plus.

Adding a vCenter Server instance

When a vCenter Server instance is added to IBM Spectrum Protect Plus, an inventory of the instance is captured, enabling you to complete backup and restore jobs, as well as run reports.

Procedure

To add a vCenter Server instance, complete the following steps:

1. In the navigation panel, expand **Manage Protection > Virtualized Systems > VMware**.
2. Click **Manage vCenter**.
3. Click **Add vCenter**.

4. Populate the fields in the **vCenter Properties** section:

Hostname/IP

Enter the resolvable IP address or a resolvable path and machine name.

Use existing user

Enable to select a previously entered user name and password for the vCenter Server instance.

Username

Enter your user name for the vCenter Server instance.

Password

Enter your password for the vCenter Server instance.

Port

Enter the communications port of the vCenter Server instance. The typical default port is 443.

Certificate

Tip:

You must use the fully qualified domain name (FQDN) to validate the VMware certificate. Using an IP address may fail to validate the VMware certificate.

In the **Certificate** field, select one of the following options to import the vCenter Server certificate to the IBM Spectrum Protect Plus server:

Use existing certificate

Click **Select** to select an existing certificate from the **Select a certificate** list.

Copy and paste

In the **Enter certificate name** field, enter a name for the certificate. Then, paste the contents of the certificate in the **Copy and paste certificate here** field and click **Create**.

Upload

- a. Download the file from the vCenter server to the local machine where you are running the browser.
- b. Click **Choose file** and search for the downloaded certificate in your system.
- c. Click **Upload**.

Get Certificate

Click **Get Certificate** to automatically retrieve the certificate that is associated with the vCenter Server that is specified in the **Hostname/IP** field.

5. In the **Options** section, configure the following option:

Maximum number of VMs to process concurrently per ESX server and per SLA

Set the maximum number of concurrent VM snapshots to process on the ESX server. The default setting is 3.

6. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the vCenter Server instance to the database, and then catalogs the instance.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

What to do next

After you add a vCenter Server instance, complete the following action:

Action	How to
Assign user permissions to the hypervisor.	See “Creating a role” on page 463 .

Related concepts

[“Managing identities” on page 469](#)

Some features in IBM Spectrum Protect Plus require credentials to access your resources. For example, IBM Spectrum Protect Plus connects to Oracle servers as the local operating system user that is specified during registration to complete tasks like cataloging, data protection, and data restore.

Related tasks

[“Backing up VMware data” on page 220](#)

Use a backup job to back up VMware resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data” on page 232](#)

VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Editing properties for a vCenter Server instance

You can edit the properties that are associated with the vCenter Server instance to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit the properties of a vCenter Server instance, complete the following steps:

1. In the navigation panel, expand **Manage Protection > Virtualized Systems > VMware**.
2. Click **Manage vCenter**.
3. Click the edit icon  that is associated with the vCenter Server instance.
4. Revise the fields in the **Edit vCenter Properties** page.

In the **Certificate** field, select one of the following options to import the vCenter Server certificate to the IBM Spectrum Protect Plus server:

Tip:

You must use the fully qualified domain name (FQDN) to validate the VMware certificate. Using an IP address may fail to validate the VMware certificate.

Use existing certificate

Click **Select** to select an existing certificate from the **Select a certificate** list.

Copy and paste

In the **Enter certificate name** field, enter a name for the certificate. Then, paste the contents of the certificate in the **Copy and paste certificate here** field and click **Create**.

Upload

- a. Download the file from the vCenter server to the local machine where you are running the browser.
- b. Click **Choose file** and search for the downloaded certificate in your system.
- c. Click **Upload**.

Get Certificate

Click **Get Certificate** to automatically retrieve the certificate that is associated with the vCenter Server that is specified in the **Hostname/IP** field.

5. Click **Save**.

vCenter Server privileges

vCenter Server privileges are required for the virtual machines that are associated with a VMware provider. These privileges are included in the vCenter Administrator role.

Virtual machine privileges

If the user that is associated with the provider is not assigned to the Administrator role for an inventory object, the user must be assigned to a role that has the following required privileges. Ensure that the privileges are propagated to child objects. For instructions for adding a permission to an inventory object, see the [Add a Permission to an Inventory Object page \(HTTPS://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-A0F6D9C2-CE72-4FE5-BAFC-309CFC519EC8.html\)](https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-A0F6D9C2-CE72-4FE5-BAFC-309CFC519EC8.html).

A test feature is available to verify that a user account has the required VMware privileges. Follow the instructions in [“Testing a vCenter Server user account for required privileges”](#) on page 218 to view the VMware privileges that are associated with the user account.

vCenter Server Object	Required Privileges
Alarm	<ul style="list-style-type: none">• Acknowledge alarm• Set alarm status
Cryptographic Operations (6.5, 6.7, 7.0, and 8.0)	<ul style="list-style-type: none">• Add disk• Direct access• Encrypt• Encrypt new• Manage encryption policies
Datastore	<ul style="list-style-type: none">• Allocate space• Browse datastore• Low level file operations• Remove datastore• Remove file• Update virtual machine files
Distributed switch	<ul style="list-style-type: none">• Port configuration operation• Port setting operation
Folder	<ul style="list-style-type: none">• Create folder
Global	<ul style="list-style-type: none">• Cancel task• Manage custom attributes• Set custom attribute
Host > Configuration	<ul style="list-style-type: none">• Storage partition configuration

vCenter Server Object	Required Privileges
vSphere Tagging (6.5, 6.7, 7.0, and 8.0)	<ul style="list-style-type: none"> • Assign or Unassign vSphere Tag • Assign or Unassign vSphere Tag on Object (7.0 and 8.0) • Create vSphere Tag • Create vSphere Tag Category • Modify UsedBy Field for Category • Modify UsedBy Field for Tag
Network	<ul style="list-style-type: none"> • Assign network
Resource	<ul style="list-style-type: none"> • Apply recommendation • Assign a vApp to resource pool • Assign virtual machine to resource pool • Migrate powered off virtual machine • Migrate powered on virtual machine • Query vMotion
Virtual Machine Change > Configuration	<ul style="list-style-type: none"> • Acquire disk lease (6.7, 7.0, and 8.0) • Add existing disk • Add new disk • Add or remove device • Advanced (6.5) • Advanced configuration (6.7, 7.0, and 8.0) • Change CPU count • Change memory (6.7, 7.0, and 8.0) • Change settings (6.7, 7.0, and 8.0) • Configure raw device (6.7, 7.0, and 8.0) • Disk change tracking (6.5) • Disk lease (6.5) • Memory (6.5) • Modify device settings • Raw device (6.5) • Reload from path • Remove disk • Rename • Settings (6.5) • Toggle disk change tracking (6.7, 7.0, and 8.0)
Virtual Machine > Guest Operations	<ul style="list-style-type: none"> • Guest Operation Modifications • Guest Operation Program Execution • Guest Operation Queries

vCenter Server Object	Required Privileges
Virtual Machine > Interaction	<ul style="list-style-type: none"> • Backup operation on virtual machine • Power Off • Power On
Virtual Machine > Inventory	<ul style="list-style-type: none"> • Register • Remove • Unregister
Virtual Machine > Provisioning	<ul style="list-style-type: none"> • Allow disk access • Allow read-only disk access • Allow virtual machine download • Allow virtual machine files upload • Mark as template • Mark as virtual machine
Virtual Machine > Snapshot management	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert to snapshot
vApp	<ul style="list-style-type: none"> • Add virtual machine • Assign resource pool • Assign vApp • Create • Delete • Power Off • Power On • Rename • Unregister • vApp resource configuration • Import

Testing a vCenter Server user account for required privileges

The user account that is associated with a vCenter Server that is registered in IBM Spectrum Protect Plus must have the required privileges to complete tasks. Use the test feature to determine that an account has the required privileges.

Before you begin

Ensure that the vCenter Server and associated user account are registered in IBM Spectrum Protect Plus. To register a vCenter Server, see [“Adding a vCenter Server instance”](#) on page 213.

Procedure

To test a vCenter Server user account for required privileges, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. Click **Manage vCenter**.

3. Locate the vCenter Server that you want to test.
4. Click the **Actions** menu that is associated with the vCenter Server, and then select **Test**. The test window opens.
5. Enter the vCenter username in domain\user format in the **Username** field and the password in the **Password** field.

Remember:

- The username and password are for an admin user for the vCenter. It may be the same or different from the user used in vCenter registration in IBM Spectrum Protect Plus.
 - The username and password entered are only being used to test the permissions of the registered user.
6. Click **Test**. The test result window opens.
 7. Verify that a green check mark is provided for each privilege. If a privilege is missing, follow the instructions for adding privileges on the [Add a Permission to an Inventory Object page \(HTTPS://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-A0F6D9C2-CE72-4FE5-BAFC-309CFC519EC8.html\)](https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-A0F6D9C2-CE72-4FE5-BAFC-309CFC519EC8.html).
 8. Click **OK** to close the test window.

Detecting VMware resources

VMware resources are automatically detected after the vCenter Server instance is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the instance was added. If a virtual machine inventory job fails, subsequent attempts to run a backup job will also fail.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. In the list of vCenters Server instances, select an instance or click the link for the instance to navigate to the resource that you want. For example, if you want to run an inventory job for an individual virtual machine in the instance, click the instance link and then select a virtual machine.
3. Click **Run Inventory**.

Testing the connection to a vCenter Server virtual machine

You can test the connection to a vCenter Server virtual machine. The test function tests the hypervisor tools. It also verifies communication with the virtual machine and tests domain name server (DNS) settings between the IBM Spectrum Protect Plus virtual appliance and the virtual machine.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. In the list of vCenters Server instances, click the link for a vCenter Server to navigate to the individual virtual machines.
3. Select a virtual machine, and then click **Select Options**.
4. Select **Catalog file metadata**.
5. For Windows-based VMs, click **Get SSL certificate thumbprint**. For Linux-based VMs, click **Get server key**.

Note: This setting will only be visible for Windows-based hosts if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences” on page 173](#).

6. Enter the username in the **Guest OS Username** and the password in the **Guest OS Password** field. Alternately, if the user already exists, select **Use existing user** and select a user in the **Select user** list.
7. Click **Test**.

Backing up VMware data

Use a backup job to back up VMware resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters with snapshots.

Before you begin

Review the following procedures and considerations before you define a backup job:

- Register the providers that you want to back up. For more instructions, see [“Adding a vCenter Server instance”](#) on page 213.
- Configure SLA policies. For more instructions, see [“Creating an SLA policy for hypervisors”](#) on page 198.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,”](#) on page 455.
- If a VM is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- If your vCenter is a VM, to help maximize data protection, have the vCenter on a dedicated datastore and backed up in a separate backup job.
- Ensure the latest version of VMware Tools is installed on VMware VMs.
- You must assign tags at the VM guest level for them to be utilized for backup exclusion rules based on tags or to be used as filtering for reports in IBM Spectrum Protect Plus.

About this task

- When backing up VMware VMs, IBM Spectrum Protect Plus downloads .vmx, .vmxf, and .nvram files if necessary, and then it transfers those files to the vSnap server as needed. For this to work successfully, the IBM Spectrum Protect Plus appliance must be able to resolve and access all protected ESXi hosts. When the appliance communicates with an ESXi host, the correct IP address must be returned.
- If a VM is protected by an SLA policy, the backups of the VM will be retained based on the retention parameters of the SLA policy, even if the VM is removed from vCenter.
- If an existing VM is migrated by a vMotion operation, IBM Spectrum Protect Plus will perform a rebase operation if necessary.

Restriction: File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a non-default local administrator is entered as the **Guest OS Username** when defining a backup job. A non-default local administrator is any user that has been created in the guest OS and has been granted the administrator role.

This occurs if the registry key `LocalAccountTokenFilterPolicy` in `[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` is set to 0 or not set. If the parameter is set to 0 or not set, a local non-default administrator cannot interact with WinRM, which is the protocol IBM Spectrum Protect Plus uses to install the Windows agent for file cataloging, send commands to this agent, and get results from it.

Set the `LocalAccountTokenFilterPolicy` registry key to 1 on the Windows guest that is being backed up with `Catalog File Metadata` enabled. If the key does not exist, navigate to `[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` and add a `DWord` Registry key named `LocalAccountTokenFilterPolicy` with a value of 1.

Procedure

To define a VMware backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. Select resources to back up.

Use the search function to search for available resources and toggle the displayed resources by using the **View** filter. Available options are **VMs and Templates**, **VMs**, **Datastore**, **Tags and Categories**, and **Hosts and Clusters**. Tags are applied in vSphere, and allow a user to assign metadata to VMs.

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

The job will run as defined by the SLA policies that you selected. To run the job immediately, click **Jobs and Operations > Schedule**. Select the job and click **Actions > Start**.

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
- To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job”](#) on page 439.

When the job definition is saved, available virtual machine disks (VMDKs) in a VM are discovered and are shown when **VMs and Templates** is selected in the **View** filter. By default, these VMDKs are assigned to the same SLA policy as the VM. If you want a more granular backup operation, you can exclude individual VMDKs from the SLA policy. For instructions, see [“Excluding VMDKs from the SLA policy for a job”](#) on page 225.

5. To edit options before you create the job definition, click **Select Options**.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

In the **Backup Options** section, set the following job definition options:

Skip Read-only datastores

Skip datastores that are mounted as read-only.

Skip temporary datastores mounted for Instant Access

Exclude temporary Instant Access datastores from the backup job definition.

VADP Proxy

Select a VADP proxy to balance the load.

Priority

Set the backup priority of the selected resource. Resources with a higher priority setting are backed up first in the job. Click the resource that you want to prioritize in the **VMware Backup** section, and then set the backup priority in the **Priority** field. Set 1 for the highest priority resource or 10 for the lowest. If a priority value is not set, a priority of 5 is set by default.

In the **Snapshot Options** section, set the following job definition options:

Make VM snapshot application/file system consistent

Enable this option to turn on application or file system consistency for the VM snapshot. All VSS-compliant applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL, and the system state are quiesced. VMDKs and VMs can be instantly mounted to restore data that is related to quiesced applications.

VM Snapshot retry attempts

Set the number of times that IBM Spectrum Protect Plus attempts to capture an application or file-consistent snapshot of a VM before the job is canceled. If the **Fall back to unquiesced snapshot if quiesced snapshot fails** option is enabled, an unquiesced snapshot will be taken after the retry attempts.

Fall back to unquiesced snapshot if quiesced snapshot fails

Enable to fall back to a non-application or non-file-system consistent snapshot if the application consistent snapshot fails. Selecting this option ensures that an unquiesced snapshot is taken if environmental issues prohibit the capture of an application or file-system consistent snapshot.

In the **Agent Options** section, set the following job definition options:

Truncate SQL logs

To truncate application logs for SQL Server databases that are on the VM during the backup job, enable the **Truncate SQL logs** option.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. Do not select this option if you want to back up the database transaction logs during the SQL Server backup operation. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

For more information about backing up SQL Server database logs, see [“Log backups” on page 395](#).

The credentials must be established for the associated VM by using the Guest OS user name and Guest OS Password option within the backup job definition. When the VM is attached to a domain, the user identity follows the default *domain\name* format. If the user is a local administrator, the format *local_administrator* is used.

The user identity must have local administrator privileges. On the SQL Server server, the system login credential must have the following permissions:

- SQL Server sysadmin permissions must be enabled.
- The **Log on as a service** right must be set. For more information about this right, see [Add the Log on as a service Right to an Account](#).

IBM Spectrum Protect Plus generates log files for the log truncation function and copies them to the following location on the IBM Spectrum Protect appliance:

```
/data/log/guestdeployer/latest_date/latest_entry/vm_name
```

where *latest_date* is the date that the backup job and log truncation occurred, *latest_entry* is the universally unique identifier (UUID) for the job, and *vm_name* is the host name or IP address of the VM where the log truncation occurred.

Restriction: File indexing and file restore are not supported from restore points that were copied to cloud resources or repository servers.

Catalog file metadata

Turn on file indexing for the associated snapshot. When file indexing is completed, individual files can be restored by using the **File Restore** pane in IBM Spectrum Protect Plus. Credentials must be established for the associated VM by using an SSH key, or the **Guest OS Username** and **Guest OS Password** options within the backup job definition. Ensure that the VM can be accessed from the IBM Spectrum Protect Plus appliance either by using DNS or a host name.

Restriction: SSH Keys are not a valid authorization mechanism for Windows platforms.

Run as system user

Run the file indexing as the system user. This allows the file indexing to be run at the highest privilege level on the client virtual machine. **Catalog file metadata** must be enabled to use this option.

Exclude Files

Enter directories to skip during file indexing. Files within these directories are not added to the IBM Spectrum Protect Plus catalog and are not available for file recovery. Directories can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *. Separate multiple filters with a semicolon. **Catalog file metadata** must be enabled to use this option.

Get SSL certificate thumbprint or Get SSH key

Note: This setting will only be visible for Windows-based hosts if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences”](#) on page 173.

Verify the identity of the VM being backed up. **Catalog file metadata** must be enabled to use this option.

For Windows-based virtual machines:

Obtain the certificate thumbprint and verify that the certificate thumbprint matches the thumbprint of the certificate on the host. Click **Get SSL certificate thumbprint**.

Get SSL certificate thumbprint

Get the SSL certificate thumbprint for the Windows-based host. You must complete this step when registering servers for the first time or if the certificate on the server changes.

The HTTPS listener must be enabled on the host. You must create a self-signed certificate and then enable the HTTPS listener if it is not already enabled. For more information, see [How to configure WinRm for HTTPS](#).

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the certificate thumbprint will automatically be added to the registration information in the catalog.

SSL certificate thumbprint

The SSL certificate thumbprint is displayed here. Confirm that the certificate thumbprint matches the thumbprint of the certificate on the host that you are adding.

For Linux-based virtual machines:

Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that they key fingerprint matches the key fingerprint of the host that you are adding.

Use existing user

Select a previously entered user name and password for the provider.

Guest OS Username/Password

For some tasks (such as cataloging file metadata, file restore, and IP reconfiguration), credentials must be established for the associated VM. Enter the user name and password, and ensure that the VM can be accessed from the IBM Spectrum Protect Plus appliance either by using DNS or a host name.

6. To troubleshoot a connection to a hypervisor VM, use the **Test** function.

The **Test** function verifies the hypervisor tool settings and tests DNS settings between the IBM Spectrum Protect Plus appliance and the VM. Select a single VM, and then click **Select Options**. You must select **Catalog file metadata**. Select **Use existing user** to select a previously entered username and password for the resource. Alternately, enter a username in the **Guest OS Username** and password in the **Guest OS Password** fields if you have not previously entered the username and password for the resource. Click **Test**. For more information, see [“Testing the connection to a vCenter Server virtual machine” on page 219](#).

7. Click **Save**.

8. To configure additional options, click the **Policy Options** clipboard icon  that is associated with the job in the **SLA Policy Status** section. Set the following additional policy options:

Pre-scripts and Post-scripts

Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and a script server where the script will run. Scripts and script servers are configured by using the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Run inventory before backup

Run an inventory job and capture the latest data of the selected resources before starting the backup job.

Exclude Resources

Exclude specific resources from the backup job by using single or multiple exclusion patterns. Resources can be excluded by using an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Exclude Resources by Tag

Exclude specific resources based on associated VM tags from the backup job. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *. Multiple filters may be separated with a semicolon.

Force Full Backup of Resources

Force base backup operations for specific VMs or databases in the backup job definition. Separate multiple resources with a semicolon.

9. To save any additional options that you configured, click **Save**.

What to do next

After you define a backup job, you can complete the following actions:

Action	How to
If you are using a Linux environment, consider creating VADP proxies to enable load sharing.	See “Creating VADP proxies” on page 227 .
Create a VMware restore job definition.	See “Restoring VMware data” on page 232 .

In some cases, VMware backup jobs fail with “failed to mount” errors. To resolve this issue, increase the maximum number of NFS mounts to at least 64 by using the NFS.MaxVolumes (vSphere 5.5 and later) and NFS41.MaxVolumes (vSphere 6.0 and later) values. Follow the instructions in [Increasing the default value that defines the maximum number of NFS mounts on an ESXi/ESX host](#).

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Starting jobs on demand” on page 433](#)

You can run any job on demand, even if the job is set to run on a schedule.

Excluding VMDKs from the SLA policy for a job

After you save a backup job definition, you can exclude individual VMDKs in a virtual machine from the SLA policy that is assigned to the job.

Before you begin

Excluding one or more VMDKs from a backup operation can impact the success of recovery. Consider the following scenarios before excluding a disk from a VM backup operation.

- For Instant Disk Restore, if a VMDK is selected for a restore operations, an existing VM is chosen as the destination. IBM Spectrum Protect Plus mounts the restored disk to the chosen destination VM.
- For Instant VM Restore, if the VMDK that was excluded during a backup contains data that is necessary to boot the virtual machine, then the restored VM may fail to boot.
- For VMs with Windows-based guests, the restored VM may fail to boot if the disk on which the main operating system is installed, typically the C : drive, was excluded during the backup operation.
- For VMs with Linux-based guests, the restored VM may fail:
 - If a disk containing the boot or root partition was excluded during backup.
 - If a disk containing a data (non-root) partition was excluded during backup, and the data volume did not have the 'nofail' option specified in /etc/fstab, then the restored VM may fail.

Procedure

To exclude VMDKs from the SLA policy that has been applied to a virtual machine:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. Select **VMs and Templates** in the **View** filter.

3. Click the link for the vCenter that contains the virtual machines in the **Name** field. It may be necessary to select the datacenter and folder.
4. Identify the virtual machine that contains the VMDKs to be excluded. If an SLA policy is already applied, it will be visible in the **SLA Policy** field.
5. Click the virtual machine name for the virtual machine to display the associated VMDKs.
6. Select the VMDK that is to be excluded and then click **Select SLA Policy**.
7. Clear the selected SLA policy that is applied to the VMDK in the SLA Policy pane.

Note: If multiple VMDKs are selected that are assigned to the same policy, the SLA policy will not appear as selected in the SLA Policy pane after clicking **Select SLA Policy**. Leave all SLA policies clear to exclude the VMDKs from the policy.

8. Click **Save**.

Managing VADP backup proxies

In IBM Spectrum Protect Plus, you must create proxies to run VMware backup jobs by using vStorage API for Data Protection (VADP) in Linux environments. The proxies reduce demand on system resources by enabling load sharing and load balancing.

The backup of a VMware virtual machine includes the following files:

- VMDKs corresponding to all disks. The base backup captures all allocated data, or all data if disks are on NFS datastores. Incremental backups will capture only changed blocks since the last successful backup.
- Virtual machine templates.
- VMware files with the following extensions:
 - .vmx
 - .vmfx (if available)
 - .nvram (stores the state of the virtual machine BIOS)

At least one VADP proxy must be enabled in the backup site that is specified in the SLA for VMware backups. For more information, see [“Creating VADP proxies” on page 227](#).

The processing load is shifted off the host system and onto the proxies for VMware backup jobs. When more than one VADP proxy exists, throttling ensures that multiple proxies are optimally utilized to maximize data throughput. For each VMware virtual machine being backed up, IBM Spectrum Protect Plus determines which VADP proxy is the least busy and has the most available memory and free tasks. Free tasks are determined by the number of available CPU cores or by using the **Softcap task limit** option.

If a proxy server goes down or is otherwise unavailable before the start of the job, the other proxies take over and the job is complete. If a proxy server becomes unavailable when a job is running, the job may fail.

Transport modes describe the method by which a VADP proxy moves data. The transport mode is set as a property of the proxy. Most backup and recovery jobs are later configured to use one or more proxies.

VADP proxies in IBM Spectrum Protect Plus support the following VMware transport modes: SAN, HotAdd, NBDSSL, and NBD.

Although every enterprise differs, and priorities in terms of size, speed, reliability, and complexity vary from environment to environment, the following general guidelines apply to the Transport Mode selection:

- SAN transport mode is preferred in a direct storage environment because this mode is typically fast and reliable.
- HotAdd transport mode is preferred if the VADP proxy is virtualized. This mode supports all vSphere storage types.

Note: To use only the HotAdd transport mode without falling back to alternate transport modes, select **VADP proxy uses only HotAdd transport mode** in **Global Preferences**. For more information, see [“Configuring global preferences” on page 173](#).

- NBD or NBDSSL transport mode (LAN) is the fallback mode because it works in physical, virtual, and mixed environments. However, with this mode, the data transfer speed might be compromised if network connections are slow. NBDSSL mode is similar to NBD mode except that data transferred between the VADP proxy and the ESXi server is encrypted when using NBDSSL.

Creating VADP proxies

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Before you begin

Review the IBM Spectrum Protect Plus system requirements in [technote 6837823](#).

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see [“Permission types” on page 463](#).

Restriction: For running the steps to create VADP proxies, ensure that you have a user ID with the SYSADMIN role assigned. For more information about roles, see [“Managing roles” on page 461](#).

Procedure

To create VMware VADP proxies, complete the following steps:

1. In the navigation panel, click **System Configuration > VADP Proxy**.
2. Click **Register Proxy**.
3. Complete the following fields in the **Install VADP Proxy** pane:

Hostname/IP

Enter the resolvable IP address or a resolvable path and machine name.

Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

Select a site

Select a site to associate with the proxy.

Use existing user

Enable to select a previously entered user name and password for the provider.

Username

Enter the user name for the VADP proxy server.

Password

Enter the password name for the VADP proxy server.

4. Click **Install**.
5. Click **Yes** on the confirm screen.
6. Repeat the previous steps for each proxy that you want to create.

Results

The proxy is added to the **VADP Proxy** table. You can suspend, uninstall, unregister, or edit a proxy server by clicking the ellipses icon ******* to open the actions menu. Suspending a proxy prevents upcoming backup jobs from using the proxy, and jobs that use a suspended or unregistered proxy will run locally, which may impact performance. You can complete maintenance tasks on the proxy while it is suspended. To resume usage of the proxy, click the ellipses icon ******* to open the actions menu and click **Resume**. After successful creation, the service `vadp` is started on the proxy machine. A log file, `vadp.log`, is generated in `/opt/IBM/SPP/logs` directory.

The connection between IBM Spectrum Protect Plus and a registered VADP proxy is a bidirectional connection that requires the IBM Spectrum Protect Plus virtual appliance to have connectivity to the VADP proxy, and the VADP proxy to have connectivity to the IBM Spectrum Protect Plus virtual appliance.

To ensure a proper connection from the IBM Spectrum Protect Plus virtual appliance to the VADP proxy, verify that IBM Spectrum Protect Plus virtual can ping the VADP proxy by completing the following steps:

1. Connect to the command line for the IBM Spectrum Protect Plus virtual appliance by using the Secure Shell (SSH) network protocol.
2. Issue the following command: `ping <vadp_ip>`, where `<vadp_ip>` is the resolvable IP address of the VADP proxy.

If the ping fails, ensure that the IP address of the VADP proxy is resolvable and is addressable by the IBM Spectrum Protect Plus appliance and that a route exists from the IBM Spectrum Protect Plus appliance to the VADP proxy.

If the ping succeeds, ensure that there is a proper connection from the VADP proxy to the IBM Spectrum Protect Plus virtual appliance by performing the following procedure:

1. Connect to the command line for the VADP proxy by using Secure Shell (SSH) network protocol.
2. Issue the following command: `ping <spectrum_protect_plus_ip>`, where `<spectrum_protect_plus_ip>` is the resolvable IP address of the IBM Spectrum Protect Plus virtual appliance.

If the ping fails, ensure that the IP address of the IBM Spectrum Protect Plus virtual appliance is resolvable and is addressable by the VADP proxy. Ensure that a route exists from the VADP proxy to the IBM Spectrum Protect Plus virtual appliance.

What to do next

After you create the VADP proxies, you can complete the following action:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 220. The proxies are indicated in the job log by a log message similar to the following text: Run <code>remote vmdkbackup of MicroService:</code> <code>http://<proxy></code> <code>nodename, IP:proxy_IP_address</code>

Related tasks

[“Setting options for VADP proxies” on page 229](#)

When you create VADP proxies in IBM Spectrum Protect Plus, you can configure various options for each VADP proxy.

Registering a VADP proxy on a vSnap server

You can install and register a VADP proxy on a physical or virtual vSnap server. When you install and register a VADP proxy locally on a vSnap server, no NFS mount is needed. Data movement is optimized because the file system is on the same machine and can be referenced directly for both backup and restore jobs. VADP proxies that are not installed and registered on a vSnap server still require an NFS mount.

Before you begin

One or more stand-alone vSnap servers must be properly deployed and configured in your environment and added to IBM Spectrum Protect Plus backup storage providers. For instructions, see [“Registering a vSnap server as a backup storage provider” on page 40](#).

For the combined system requirements of a vSnap server and the VADP proxy, see [VADP proxy on vSnap server requirements](#).

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see [“Permission types” on page 463](#).

The identity associated with a vSnap server is the account that is used to register the VADP proxy on the vSnap server. When you register a VADP proxy on a vSnap server, an installer is pushed and requires sudo privileges to successfully install the VADP proxy software. The identify associated with a vSnap server must have sudo privileges.

Tip: Use the `serveradmin` User ID when adding a vSnap server to IBM Spectrum Protect Plus. When you deploy a VADP proxy to a vSnap server, this account is used which already has all of the necessary privileges.

Procedure

1. In the navigation panel, click **System Configuration > Storage > vSnap servers**.
2. Select the vSnap server on which the VADP proxy is to be installed and registered.
3. Click **Register as VADP Proxy**.
4. In the Confirm dialog box, click **Yes**.

Results

When the process is complete, a green checkmark will appear in the **VADP Proxy** column in the table of the Disk Storage pane.

Setting options for VADP proxies

When you create VADP proxies in IBM Spectrum Protect Plus, you can configure various options for each VADP proxy.

Before you begin

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see [“Permission types” on page 463](#).

Settings for VADP proxies are not saved until you save them by clicking the **Save** button.

Procedure

To set options for VMware VADP proxies, complete the following steps:

1. In the navigation panel, click **System Configuration > VADP Proxy**.
2. Click the VADP proxy that you want to configure, which then displays the information in the adjacent details pane.
3. In the VADP proxy details pane, click the ellipses icon ******* and then choose **Proxy Options**.
4. Complete the following fields in the **Set VADP Proxy Options** pane:

Transport Modes (ordered list)

Set the transport modes to be used by the proxy. The listed transport mode is an ordered list to ensure that the optimal order is select. Choose one of the following options: **san:hotadd:nbdssl:nbd**, **san:hotadd:nbd**, **san:hotadd**, **san:nbdssl:nbd**, **san:nbd**, **san**, **hotadd:nbdssl:nbd**, **hotadd:nbd**, **hotadd**, **nbdssl:nbd**, **nbdssl**, and **nbd**. The order in which each mode is listed will determine the order in which the transport modes are used. For more information about VMware transport modes, see [Virtual Disk Transport Methods](#).

Enable NBDSSL Compression

If you selected the NBDSSL transport mode, enable compression to increase the performance of data transfers. Available compression types include **libz**, **fastlz**, and **skipz**.

To turn off compression, select **disabled**.

Log retention in days

Set the number of days to retain logs before they are deleted. Available options are **7 Days**, **14 Days**, **30 Days**, and **60 Days**.

Read and write buffer size

Set the buffer size of the data transfer, measured in bytes. Available options are **64K**, **128K**, **256K**, **512K**, **1024K**, **2048K**, and **4096K**.

Block size of NFS volume

Set the block size to be used by the mounted NFS volume, measured in bytes. Available options are **64K**, **128K**, **256K**, **512K**, and **1024K**.

Softcap task limit

Set the number of concurrent VMs that a proxy can process. If **Use All Resources** is selected, the number of CPUs on the proxy determines the task limit based on the following formula:

$$1 \text{ CPU} = 1 \text{ VMDK}$$

A CPU is the smallest hardware unit capable of executing a thread. The number of CPUs on a proxy is determined by using the `lscpu` command.

Other available options are **4**, **8**, **12**, **16**, **20**, **24**, **28**, **32**, and **36**.

What to do next

After setting the VADP proxy options, you can complete the following actions:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 220.
Uninstall the proxies when you cease running the VMware backup jobs.	See “Uninstalling VADP proxies” on page 231.

Related tasks

[“Creating VADP proxies”](#) on page 227

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Uninstalling VADP proxies

You can remove a VADP proxies from your IBM Spectrum Protect Plus environment.

Procedure

To uninstall VADP proxies from your IBM Spectrum Protect Plus, complete the following steps:

Note: This procedure only applies to VADP proxies that have been installed in the environment. It does not apply to the VADP proxy that is automatically deployed on IBM Spectrum Protect Plus in previous versions.

1. In the navigation panel, click on **System Configuration > VADP Proxy**.
2. Click the VADP proxy that you want to uninstall, which then displays the information in the adjacent details pane.
3. Click the ellipses icon **⋮** in the details pane and select **Uninstall**.

Updating the IBM Spectrum Protect Plus IP address for VADP proxies

You can use the IBM Spectrum Protect Plus user interface to deploy vStorage API for Data Protection (VADP) proxy servers. If the IP address of the IBM Spectrum Protect Plus server changes after VADP proxy deployment, the proxies in the environment lose contact with the IBM Spectrum Protect Plus server. To resolve this issue, a script is provided to update the IBM Spectrum Protect Plus address for orphaned VADP proxies.

Before you begin

About this task

Ensure that you have the new IP address of the IBM Spectrum Protect Plus server.

Procedure

1. If the VADP proxy server and vSnap server are co-deployed, log in by using the secure shell protocol (SSH) with the `serveradmin` user account. If you are using a stand-alone VADP proxy server, log in to the VADP proxy server with the `root` user account.
2. Run the `update_vadp.sh` script with the new IBM Spectrum Protect Plus server IP address as the only argument by taking one of the following actions:
 - If the VADP proxy is co-deployed with the vSnap server, run the following script as the `serveradmin` user:

```
$ sudo /opt/IBM/SPP/bin/update_vadp.sh new_ip
```

- If the VADP proxy server is stand-alone, run the script as the `root` user:

```
# /opt/IBM/SPP/bin/update_vadp.sh new_ip
```

3. When the script completes processing, the following message is displayed:

```
The new IP address ( <new_ip> ) has been successfully updated.
```

The IBM Spectrum Protect Plus server IP address is updated on the VADP proxy server

Restoring VMware data

VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Before you begin

Complete the following tasks:

- Ensure that a VMware backup job was run at least once. For instructions, see [“Backing up VMware data” on page 220](#).
- Ensure that appropriate roles are assigned to IBM Spectrum Protect Plus users so that they can complete backup and restore operations. Grant users access to hypervisors and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,” on page 455](#) and [“Managing user accounts” on page 466](#).
- Ensure that the destination that you plan to use for the restore job is registered in IBM Spectrum Protect Plus. This requirement applies to restore jobs that restore data to original hosts or clusters.
- When restoring a virtual machine by using clone mode and by using the original IP configuration, ensure that credentials are established through the **Guest OS Username** and **Guest OS Password** options within the backup job definition.
- You must assign tags at the VM guest level for them to be utilized for backup exclusion rules based on tags or to be used as filtering for reports in IBM Spectrum Protect Plus.

About this task

If a VMDK is selected for restore operation, IBM Spectrum Protect Plus automatically presents options for an Instant Disk restore job, which provides instant writable access to data and application restore points. An IBM Spectrum Protect Plus snapshot is mapped to a target server where it can be accessed or copied as required.

All other sources are restored through Instant VM restore jobs, which can be run in the following modes:

Production mode

Production mode enables disaster recovery at the local site from primary storage or a remote disaster recovery site, replacing original machine images with recovery images. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data jobs associated with the virtual machine continue to run. As part of a production mode restore, you can opt to replace the storage in the virtual machine with a virtual disk from a previous virtual machine backup. Select

Test mode

Test mode creates temporary virtual machines for development or testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. Test machines are kept running as long as needed to complete testing and verification and are then cleaned up. Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines used for production. Virtual machines that are created in test mode are also given unique names and identifiers to avoid conflicts within your production environment.

Clone mode

Clone mode creates copies of virtual machines for use cases that require permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines created in clone mode are also given unique names and identifiers to avoid conflicts within your production environment. With clone mode, you must be sensitive to resource consumption because clone mode creates permanent or long-term virtual machines.

Note: Restoring a VM in test mode or clone mode without selecting a fenced network may interfere with production VMs and network services. For instructions for creating a fenced network, see [“Creating a fenced network through a VMware restore job” on page 239](#).

The size of a virtual machine that is restored from a vSnap copy to an IBM Spectrum Protect restore point will be equal to the thick provisioned size of the virtual machine, regardless of source provisioning due to

the use of NFS datastores during the copy operation. The full size of the data must be transferred even if it is unallocated in the source virtual machine.

When you restore VMware data from an IBM Spectrum Protect archive, files initially will be migrated from tape to a staging pool. Depending on the size of the restore operation, this process could take several hours.

Restriction: Windows file indexing and file restore on volumes residing on dynamic disks is not supported.

Procedure

To define a VMware restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > VMware**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:

- a) Review the available sources, including virtual machines (VMs) and virtual disks (VDisks). Use the **View** filter to toggle the displayed sources to show hosts and clusters, VMs, or tags and categories. You can expand a source by clicking its name.

You can also enter all or part of a name in the **Search for** box to locate VMs that match the search criteria. You can use the wildcard character (*) to represent all or part of a name. For example, vm2* represents all resources that begin with "vm2".

- b) Click the plus icon  next to the item that you want to add to the restore list next to the list of sources. You can add more than one item of the same type (VM or virtual disk).

To remove an item from the restore list, click the minus icon  next to the item.

- c) Click **Next**.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:

Option	Description
	<ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resource restore or recurring restore

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>

Option	Description
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Demo The demonstration site from which to restore snapshots. This menu item is available only if you updated the product from IBM Spectrum Protect Plus 10.1.6 or earlier.</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

- On the **Set destination** page, specify the instance that you would like to restore for each chosen source and click **Next**:

Original Host or Cluster

Select this option to restore data to the original host or cluster.

Alternate Host or Cluster

Select this option to restore data to a local destination that is different from the original host or cluster, and then select the alternate location from the available resources. Test and production networks can be configured on the alternate location to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. From the **vCenters** section, select an alternative location. You can filter the alternative locations by either hosts or clusters.

In the **VM Folder Destination** field, enter the virtual machine folder path on the destination datastore. Note that the directory will be created if it does not exist. Use "/" as the root virtual machine folder of the targeted datastore.

When restoring a virtual disk to a new destination VM, select the virtual machine to which the virtual disk will be restored and the **Destination Disk Mode**. You can set the **Destination Controller** to select a supported SCSI controller. Changing the SCSI controller type replaces the existing controller with a new controller, applies the common settings of the existing controller to the new controller, and reassigns all SCSI devices to the new controller. Optionally, can also set the **Destination Controller Address #** and **Destination Controller LUN #** to select specific controllers or LUNs.

ESX host if vCenter is down

Select this option to bypass vCenter Server and to restore data directly to an ESXi host. In other restore scenarios, actions are completed through vCenter Server. If vCenter Server is unavailable,

this option restores the virtual machine or virtual machines that contain the components that vCenter Server is dependent on.

When you select an ESXi host, you must specify the host user. You can select an existing user for the host or create a new one.

To create a user, enter a user name, the user ID, and the user password.

If the ESXi host is attached to a domain, the user ID follows the default *domain\name* format. If the user is a local administrator, use the *local_administrator* format.

To restore data to an ESXi host, the host must have a standard switch or a distributed switch with ephemeral binding. Review the information in [“Restoring data when vCenter Server or other management VMs are not accessible” on page 240](#) to ensure that you have the correct environment configured to use this option.

6. On the **Set datastore** page, take the following actions:

- If you are restoring data to an alternate ESXi host or cluster, select the destination datastore and click **Next**.
- If you are restoring data to the original ESXi host or cluster, this page is not displayed.

7. On the **Set network** page, specify the network settings to use for each chosen source and click **Next**.

- If you are restoring data to the original ESXi host or cluster, specify the following network settings:

Allow OS to define IP configuration

Select this option to allow your operating system to define the destination IP address. During a test mode restore operation, the destination virtual machine receives a new MAC address along with an associated NIC. Depending on your operating system, a new IP address can be assigned based on the original NIC of the virtual machine, or assigned through DHCP. During a production mode restore, the MAC address does not change; therefore, the IP address should be retained.

Use original IP configuration

Select this option to restore data to the original host or cluster using your predefined IP address configuration. During the restore operation, the destination virtual machine receives a new MAC address, but the IP address is retained.

- If you are restoring data to an alternate ESXi host or cluster, complete the following steps:
 - a. In the **Production** and **Test** fields, set virtual networks for production and test restore job runs. Destination network settings for production and test environments should point to different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks that are associated with test and production modes will be used when the restore job is run in the associated mode.
 - b. Set an IP address or subnet mask for virtual machines to be repurposed for development, testing, or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines that contain multiple NICs are supported.

Take one of the following actions:

- To allow your operating system to define the destination subnets and IP addresses, click **Use system defined subnets and IP addresses for VM guest OS on destination**.
- To use your predefined subnets and IP addresses, click **Use original subnets and IP addresses for VM guest OS on destination**.
- To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, click **Add Mapping**, and enter a subnet or IP address in the **Add Source Subnet or IP Address** field.

Choose one of the following network protocols:

- Select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected source.

- Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. The **Subnet / IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

Note: When a mapping is added, the source IP address must be entered into the field by the **+** button. The destination IP address information should be entered into the **Subnet / IP Address**, **Subnet Mask**, and **Gateway** fields. Re-addressing can only be performed on machines with VMware Tools installed prior to executing the backup job that is to be restored.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source virtual machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine uses DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment, all addresses are assumed to be static, and only IP mapping will be available.

8. On the **Restore methods** page, select the restore method to be used for source selection. Set the VMware restore job to run in **Production**, **Test**, or **Clone**. When you run in production mode, you have the option to select **Replace virtual disks only and retain virtual machine configuration**. Enabling this option replaces the storage in the virtual machine with the virtual disks from a previous virtual machine backup. This restore method maintains the virtual machine configuration replacing only the storage. When production restore is selected, the virtual machine to which the disk replace is applied must be powered off and only restoring to the original location is supported. Additionally, options such as overwriting the virtual machine and restoring based on tags are not available because the virtual machine is not being recreated. After the job is created, it can be run in production or clone mode through the **Job Sessions** pane.
9. You can also change the name of the restored VM by entering the new VM name in the **Rename VM (optional)** field. The **Rename VM** option is not available when the production mode is used with the replace virtual disks option enabled. Click **Next** to continue.
10. If you are running the restore job in advanced mode, you can set additional options as follows:

Power® on after recovery

Toggle the power state of a virtual machine after a recovery is run. Virtual machines are powered on in the order in which they are recovered, as set in the Source step. If **Use original IP configuration** is selected, the **Power on after recovery** option is not honored.

Restriction: Restored virtual machine templates cannot be powered on after recovery.

Overwrite virtual machine

Enable this option to allow the restore job to overwrite the selected virtual machine. By default, this option is disabled.

Continue with restore even if it fails

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the restore job stops if the recovery of a resource fails.

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow to overwrite and force cleanup of pending old sessions

Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Restore VM tags

Enable this option to restore tags that are applied to virtual machines through vSphere.

Enable Streaming (VADP) restore

Parallel streaming for virtual machine restore operations is set by default. You can deselect this option for virtual machine restore operations. VADP proxy selection for streaming restore is based on the site of the source snapshot being restored.

Tip: When you are restoring virtual machines managed by a VMware Cloud (VMC) on AWS Software-Defined Data Center (SDDC), this option should always be enabled to allow streaming of the data.

Append suffix to virtual machine name

Enter a suffix to add to the names of restored virtual machines.

Prepend prefix to virtual machine name

Enter a prefix to add to the names of restored virtual machines.

11. Optional: On the **Apply scripts** page, choose the following script options and click **Next**.

- Select **Pre-script** to select an uploaded script, and an application or script server where the prescript runs. To select an application server where the script will run, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
- Select **Post-script** to select an uploaded script and an application or script server where the postscript runs. To select an application server where the script runs, clear the **Use Script Server** check box. Navigate to the **System Configuration > Script** page to configure scripts and script servers.
- Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns COMPLETED. If a postscript completes with a nonzero return code, the postscript task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a FAILED status.

12. Take one of the following actions on the **Schedule** page:

- To run an on-demand job, click **Next**.
- To set up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.

13. On the **Review** page, review your restore job settings and click **Submit** to create the job.

On-demand jobs will begin immediately; recurring jobs will begin at the scheduled start time.

What to do next

After the job is completed, select one of the following options from the **Actions** menu on the Jobs Sessions or Active Clones sections in the **Restore** pane:

Cleanup

Destroys the virtual machine and cleans up all associated resources. Because this is a temporary virtual machine to be used for testing, all data is lost when the virtual machine is destroyed.

Move to Production (vMotion)

Migrates the virtual machine through vMotion to the datastore and the virtual Network defined as the production network.

Clone (vMotion)

Migrates the virtual machine through vMotion to the datastore and virtual Network defined as the test network.

Related tasks

[“Adding a vCenter Server instance” on page 213](#)

When a vCenter Server instance is added to IBM Spectrum Protect Plus, an inventory of the instance is captured, enabling you to complete backup and restore jobs, as well as run reports.

Creating a fenced network through a VMware restore job

Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines that are used for production. Fenced networking can be used with jobs that are running in test mode and production mode.

Before you begin

Create and run a VMware Restore job. For instructions, see [“Restoring VMware data” on page 232](#).

Procedure

To create a fenced network, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > VMware**.
2. In the **Restore** pane, review the available restore points of your VMware sources, including virtual machines, VM templates, datastores, folders, and vApps. Use the search function and filters to fine-tune your selection across specific recovery site types. Expand an entry in the **Restore** pane to view individual restore points by date.
3. Select restore points and click the add to restore list icon  to add the restore point to the Restore List. Click the remove icon  to remove items from the Restore List.
4. Click **Options** to set the job definition options.
5. Select **Alternate ESX Host or Cluster**, then select an alternate host or cluster from the vCenter list.
6. Expand the **Network Settings** section. From the **Production** and **Test** fields, set virtual networks for production and test Restore job runs. Destination network settings for production and test environments should be different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks associated with Test and Production will be utilized when the restore job is run in the associated mode. The IP addresses of the target machine can be configured by using the following options:

Use system defined subnets and IP addresses for VM guest OS on destination

Select to allow your operating system to define the destination IP address. During a Test Mode restore, the destination virtual machine receives a new MAC address along with an associated NIC. Depending on your operating system, a new IP address can be assigned based on the original NIC of the virtual machine, or assigned through DHCP. During a Production Mode restore operation the MAC address does not change; therefore, the IP address should be retained.

Use original subnets and IP addresses for VM guest OS on destination

Select to restore to the original host or cluster using your predefined IP address configuration. During a restore, the destination virtual machine receives a new MAC address, but the IP address is retained.

Set the network settings for a restore to an alternate or long distance ESX host or cluster:

From the **Production** and **Test** fields, set virtual networks for production and test restore job runs. Destination network settings for production and test environments should be different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks associated with Test and Production will be utilized when the restore job is run in the associated mode.

Set an IP address or subnet mask for virtual machines to be re-purposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported.

By default, the **Use system defined subnets and IP addresses for VM guest OS on destination** option is enabled. To use your predefined subnets and IP addresses, select **Use original subnets and IP addresses for VM guest OS on destination**.

To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, then click **Add Mapping**. Enter a subnet or IP address in the **Source** field. In the destination field, select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected client. Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. Note that **Subnet / IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine is DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment all addresses are assumed to be static, and only IP mapping will be available.

Destination Datastore

Set the destination datastore for a restore to an alternate ESX host or cluster.

VM Folder Destination

Enter the VM folder path on the destination datastore. Note that the directory will be created if it does not exist. Use "/" as the root VM folder of the targeted datastore.

7. Click **Save** to save the policy options.
8. After the job is complete, select one of the following options from the **Actions** menu on the Jobs Sessions or Active Clones sections on the **Restore** pane:

Cleanup

Destroys the virtual machine and cleans up all associated resources. Since this is a temporary/testing virtual machine, all data is lost when the virtual machine is destroyed.

Move to Production (vMotion)

Migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the "Production" Network.

Clone (vMotion)

Migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "Test" network.

Related tasks

[“Adding a vCenter Server instance” on page 213](#)

When a vCenter Server instance is added to IBM Spectrum Protect Plus, an inventory of the instance is captured, enabling you to complete backup and restore jobs, as well as run reports.

Restoring data when vCenter Server or other management VMs are not accessible

IBM Spectrum Protect Plus provides an option to automatically restore data by using an ESXi host if vCenter Server or one of the components that it uses are not accessible. This option restores the virtual machines that contain the components that vCenter Server uses.

Before you begin

To complete this procedure, you must be familiar with the ESXi and vCenter Server user interfaces.

About this task

vCenter Server uses the following components:

- Platform Services Controller (PSC)
- Software-Defined Data Center (SDDC)
- Active Directory (AD)
- Domain Name System (DNS) servers

To use the **ESX host if vCenter is down** option, the ESXi host must have a standard switch or a distributed switch. The distributed switch must have ephemeral binding. If one or both of these switches are available, you can run a restore operation in IBM Spectrum Protect Plus with the option enabled as described in [“Restoring VMware data” on page 232](#) and no further manual configuration is required.

If neither of these switches is available, you must complete the following steps before you can use the **ESX host if vCenter is down** option.

Procedure

1. Connect to the destination ESXi host user interface and create a standard virtual switch.
The new switch has no port groups or uplinks.
2. Use the Secure Shell (SSH) protocol to connect to the ESXi host.
3. List the distributed switches that are configured on the ESXi host by issuing the following command:

```
#esxcli network vswitch dvs vmware list
```

4. Identify the physical network interface card (NIC) and the port group of the distributed switch that you want to use for the restore operation.
5. Remove the physical NIC and port group from the distributed switch by issuing the following command:

```
#esxcfg-vswitch -Q physical_vnic -V port_group switch_name
```

6. Add the physical NIC and port group to the new standard switch by issuing the following command:

```
#esxcli network vswitch standard uplink add --uplink-name=physical_vnic --vswitch-name=new_standard_vswitch
```

7. In the ESXi host user interface, add a temporary port group and select the standard switch that you created in step [“1” on page 241](#).
The standard switch has one port group and one uplink.
8. Run a restore operation in IBM Spectrum Protect Plus with the **ESX host if vCenter is down** option enabled.
For instructions about running a restore operation, see [“Restoring VMware data” on page 232](#).
9. In the ESXi host user interface for the ESXi host, power on the VMs that are restored.
10. Log in to the vCenter Server user interface and start the migration of the management VMs from the temporary port group that you created in step [“7” on page 241](#) to an available distributed port group.
11. After all of the VMs are migrated to the original port group, reincorporate the physical NIC and the port group into the original distributed switch by taking the following actions. For example purposes, the following commands reference a virtualized Network Interface Card (VNIC) named `vmnic0` that is part of port group 64.

- a. Remove the network cards (known as `vmnics`) from a standard switch by issuing the following command:

```
#esxcli network vswitch standard uplink remove --uplink-name=vmnic --vswitch-name=vSwitch
```

For example:

```
#esxcli network vswitch standard uplink remove --uplink-name=vmnic0 --vswitch-name=vered_recovery
```

- b. Add network cards to the distributed switch by issuing the following command:

```
#esxcfg-vswitch -P vmnic -V unused_distributed_switch_port_ID distributed_switch
```

For example:

```
#esxcfg-vswitch -P vmnic0 -V 64 SDDC-Dswitch-Private
```

12. Delete the temporary port group and the standard switch from the ESXi host user interface.
13. After the VMs are migrated and accessible, use the ESXi host user interface to unregister, but not delete, the old VMs if the original host is reachable.

By using this method, you avoid creating duplicated information such as names, Media Access Control (MAC) addresses, operating system level IDs, and VM Universal Unique Identifiers (UUIDs). You must complete this step even if you are using a new datastore.

In some vSphere or ESXi versions, the unregister operation can be completed by using the **Remove from inventory** option. This option unregisters a VM from the vCenter Server catalog, but leaves VMDK files on the datastore where the files consume storage space. After you have fully recovered the VM and the environment is successfully running, you can regain the space by manually removing these files from the datastore.

Enabling transport encryption for VMware data

You can enable transport encryption on VMware to protect VMware data.

IBM Spectrum Protect Plus 10.1.13 introduces **Transport encryption** to protect VMware data. You can protect the data transport between the vSnap and a remote VADP by enabling **Transport encryption**. If the VADP is running on the vSnap, that path is always protected because it is a local file system access.

The VADP on the Open Snap Store Manager (OSSM) does not have **Transport encryption**. OSSM does not support remote VADP.

Review the following considerations and options:

VM Backup

When you backup VMware data, the VADP reads the data from the data store and sends it to vSnap. IBM Spectrum Protect Plus transport encryption does not apply to the data store connection. To use encrypted network-based transport for the path between the data store and the VADP proxy, the user must use the **Transport Mode**. The transport mode is defined in the **System Configuration > VADP > Proxy Options > Transport Modes**. The VMware transport modes such as SAN, HotAdd, and NBDSSL are considered secure. NBD transport mode does not support transport encryption.

Note: Ensure that the network between the data store (TODO or ESX) and VADP is secure.

VMware streaming restore

The same recommendations apply as for backups. The Streaming restore is the default configuration for VMware production and clone restore operations.

VMware non-streaming restore

Non-streaming restore operations use NFS datastore mounts for instant disk access, VM file restore, and test restore. NFS can only be secure in a separate, non-routable network. Either physically or through VLAN tagging.

Backing up and restoring Hyper-V data

To protect Hyper-V data, first add Hyper-V servers in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for the content of the servers.

Ensure that your Hyper-V environment meets the system requirements in [“Hypervisor \(Microsoft Hyper-V and VMware\) and cloud instance \(Amazon EC2\) backup and restore requirements ”](#) on page 21.

Adding a Hyper-V server

When a Hyper-V server is added to IBM Spectrum Protect Plus, an inventory of the server is captured, enabling you to complete backup and restore jobs, as well as run reports.

Before you begin

Note the following considerations and procedures before adding a Hyper-V server to IBM Spectrum Protect Plus:

- Hyper-V servers can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Protect Plus. If the Hyper-V server is part of a cluster, all nodes in the cluster must be resolvable through DNS. If DNS is not available, the server must be added to the `/etc/hosts` file on the IBM Spectrum Protect Plus appliance. If more than one Hyper-V server is set up in a cluster environment, all of the servers must be added to `/etc/hosts`. When registering the cluster in IBM Spectrum Protect Plus, register the Failover Cluster Manager.
- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI initiator Service running in their Services list. Set the service to Automatic so that it is available when the machine boots.
- Add the user to the local administrator group on the Hyper-V server.

Procedure

To add a Hyper-V server, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Hyper-V**.
2. Click **Manage Hyper-V Server**.
3. Click **Add Hyper-V Server**.
4. Populate the fields in the **Server Properties** pane:

Hostname/IP

Enter the resolvable IP address or a resolvable path and machine name.

Use existing user

Enable to select a previously entered user name and password for the server.

Username

Enter your user name for the server.

Password

Enter your password for the server.

Port

Enter the communications port of the server you are adding. The typical default port is 5985.

Select the **Use TLS** check box to enable an encrypted Transport Layer Security (TLS) connection.

If you do not select **Use TLS**, you must complete additional steps on the Hyper-V server. See [“Enabling WinRM for connection to Hyper-V servers” on page 244](#).

5. In the **Options** section, configure the following option:

Maximum number of VMs to process concurrently per Hyper-V server

Set the maximum number of concurrent virtual machine snapshots to process on the Hyper-V server.

6. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the server to the database, and then catalogs the server.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

What to do next

After you add the Hyper-V server, complete the following action:

Action	How to
Assign user permissions to the hypervisor.	See “Creating a role” on page 463.

Related tasks

“Backing up Hyper-V data” on page 245

Use a backup job to back up Hyper-V data with snapshots.

“Restoring Hyper-V data” on page 251

Hyper-V restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Enabling WinRM for connection to Hyper-V servers

If you cannot use TLS to enable encrypted network traffic between IBM Spectrum Protect Plus Hyper-V servers, you must configure WinRM on the host to allow unencrypted network traffic. Ensure that you understand the security risks that are associated with allowing unencrypted network traffic.

Procedure

To configure WinRM for connection to Hyper-V hosts:

1. On the Hyper-V host system, log in with an administrator account.
2. Open a Windows command prompt. If User Account Control (UAC) is enabled, you must open the command prompt with elevated privileges by running with the **Run as administrator** option enabled.
3. Enter the following command to configure WinRM to allow unencrypted network traffic:

```
winrm s winrm/config/service @{AllowUnencrypted="true"}
```

4. Verify that the AllowUnencrypted option is set to true through the following command:

```
winrm g winrm/config/service
```

Detecting Hyper-V resources

Hyper-V resources are automatically detected after the Hyper-V server is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the server was added. If a virtual machine inventory job fails, subsequent attempts to run a backup job will also fail.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Hyper-V**.
2. In the list of Hyper-V servers, select a server or click the link for the server to navigate to the resource that you want. For example, if you want to run an inventory job for an individual virtual machine in a server, click the server link and then select a virtual machine.
3. Click **Run Inventory**.

Testing the connection to a Hyper-V Server virtual machine

You can test the connection to a Hyper-V Server virtual machine. The test function tests the hypervisor tools. It also verifies communication with the virtual machine and tests DNS settings between the IBM Spectrum Protect Plus virtual appliance and the virtual machine.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Hyper-V**.

2. In the list of Hyper-V Servers, click the link for a Hyper-V Server virtual machine to navigate to the individual virtual machines.
3. Select a virtual machine, and then click **Select Options**.
4. Select **Catalog file metadata**.
5. For Windows-based VMs, click **Get SSL certificate thumbprint**. For Linux-based VMs, click **Get server key**.

Note: This setting will only be visible for Windows-based hosts if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences”](#) on page 173.
6. Enter the username in the **Guest OS Username** and the password in the **Guest OS Password** field. Alternately, if the user already exists, select **Use existing user** and select a user in the **Select user** list.
7. Click **Test**.

Backing up Hyper-V data

Use a backup job to back up Hyper-V data with snapshots.

Before you begin

Review the following procedures and considerations before you define a backup job:

- Register the providers that you want to back up. For more information see [“Adding a Hyper-V server”](#) on page 243
- Configure SLA policies. For instructions, see [“Creating an SLA policy for hypervisors”](#) on page 198.
- Hyper-V Backup and Restore jobs require the installation of the latest Hyper-V integration services.

For Microsoft Windows environments, see [Supported Windows guest operating systems for Hyper-V on Windows Server](#).

For Linux environments, see [Supported Linux and FreeBSD virtual machines for Hyper-V on Windows](#).

- IBM Spectrum Protect Plus uses Resilient Change Tracking (RCT) for tracking the changed blocks of Hyper-V virtual machine disks. For more information, see [Resilient Change Tracking](#).
- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI initiator Service running in their Services list. Set the service to Automatic so that it is available when the machine boots.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,”](#) on page 455.
- If a virtual machine (VM) is associated with multiple SLA Policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA Policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- If the IP address of the IBM Spectrum Protect Plus appliance is changed after an initial Hyper-V base backup is created, the target IQN of the Hyper-V resource may be left in a bad state. To correct this issue, from the Microsoft iSCSI Initiator tool, click the **Discovery** tab. Select the old IP address, then click **Remove**. Click the **Target** tab and disconnect the reconnecting session.
- If a VM is protected by an SLA policy, the backups of the VM will be retained based on the retention parameters of the SLA policy, even if the VM is removed.

About this task

Restriction: File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a non-default local administrator is entered as the **Guest OS Username** when defining a backup job. A non-default local administrator is any user that has been created in the guest OS and has been granted the administrator role.

This occurs if the registry key `LocalAccountTokenFilterPolicy` in `[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` is set to 0 or not set. If the parameter is set to 0 or not set, a local non-default administrator cannot interact with WinRM, which is the protocol IBM Spectrum Protect Plus uses to install the Windows agent for file cataloging, send commands to this agent, and get results from it.

Set the `LocalAccountTokenFilterPolicy` registry key to 1 on the Windows guest that is being backed up with Catalog File Metadata enabled. If the key does not exist, navigate to `[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` and add a DWord Registry key named `LocalAccountTokenFilterPolicy` with a value of 1.

Procedure

To define a Hyper-V backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Hyper-V**.

2. Select resources to back up.

Use the search function to search for available resources and toggle the displayed resources through the **View** filter. Available options are **VMs** and **Datastore**.

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup criteria to the job definition.

4. To create the job definition by using default options, click **Save**.

The job runs as defined by the SLA policies that you selected. To run the job manually, click **Jobs and Operations > Schedule**. Select the job and click **Actions > Start**.

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
- To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job”](#) on page 439.

5. To edit options before you start the job, click the edit icon in the table **Select Options**.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

In the **Backup Options** section, set the following job definition options:

Skip Read-only datastores

Enable to skip datastores mounted as read-only.

Skip temporary datastores mounted for Instant Access

Enable to exclude temporary Instant Access datastores from the backup job definition.

Priority

Set the backup priority of the selected resource. Resources with a higher priority setting are backed up first in the job. Click the resource that you want to prioritize in the **Hyper-V Backup** section, and then

set the backup priority in the **Priority** field. Set 1 for the highest priority resource or 10 for the lowest. If a priority value is not set, a priority of 5 is set by default.

In the **Snapshot Options** section, set the following job definition options:

Make VM snapshot application/file system consistent

Enable this option to turn on application or filesystem consistency for the VM snapshot.

VM Snapshot retry attempts

Set the number of times IBM Spectrum Protect Plus should attempt to snapshot a VM before canceling the job.

In the **Agent Options** section, set the following job definition options:

Truncate SQL logs

To truncate application logs for SQL during the Backup job, enable the **Truncate SQL logs** option.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. Do not select this option if you want to back up the database transaction logs during the SQL Server backup operation. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

For more information about backing up SQL Server database logs, see [“Log backups” on page 395](#).

The credentials must be established for the associated VM through the Guest OS Username and Guest OS Password option within the backup job definition. The user identity follows the default *domain\name* format if the VM is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

The user identity must have local administrator privileges. Additionally, on the SQL server, the system login credential must have SQL sysadmin permissions enabled, as well as the **Log on as a service** right. For more information about this right, see [Add the Log on as a service Right to an Account](#).

IBM Spectrum Protect Plus generates logs pertaining to the log truncation function and copies them to the following location on the IBM Spectrum Protect Plus appliance:

```
/data/log/guestdeployer/latest_date/latest_entry/vm_name
```

Where *latest_date* is the date that the backup job and log truncation occurred, *latest_entry* is the universally unique identifier (UUID) for the job, and *vm_name* is the hostname or IP address of the VM where the log truncation occurred.

Restriction: File indexing and file restore are not supported from restore points that were copied to an IBM Spectrum Protect server.

Catalog file metadata

To turn on file indexing for the associated snapshot, enable the Catalog file metadata option. After file indexing is complete, individual files can be restored by using the **File Restore** pane in IBM Spectrum Protect Plus. Note that credentials must be established for the associated VM by using an SSH key, or a Guest OS Username and Guest OS Password option in the backup job definition. Ensure that the VM can be accessed from the IBM Spectrum Protect Plus appliance either by using DNS or hostname. Note that SSH keys are not a valid authorization mechanism for Windows platforms.

Run as system user

Run the file indexing as the system user. This allows the file indexing to be run at the highest privilege level on the client virtual machine. **Catalog file metadata** must be enabled to use this option.

Exclude Files

Enter directories to skip when file indexing is performed. Files within these directories are not added to the IBM Spectrum Protect Plus catalog and are not available for file recovery. Directories can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support

standard alphanumeric characters as well as the following special characters: - _ and *. Separate multiple filters with a semicolon. **Catalog file metadata** must be enabled to use this option.

Get SSL certificate thumbprint or Get SSH key

Note: This setting will only be visible for Windows-based hosts if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences”](#) on page 173.

Verify the identity of the VM being backed up. **Catalog file metadata** must be enabled to use this option.

For Windows-based virtual machines:

Obtain the certificate thumbprint and verify that the certificate thumbprint matches the thumbprint of the certificate on the host. Click **Get SSL certificate thumbprint**.

Get SSL certificate thumbprint

Get the SSL certificate thumbprint for the Windows-based host. You must complete this step when registering servers for the first time or if the certificate on the server changes.

The HTTPS listener must be enabled on the host. You must create a self-signed certificate and then enable the HTTPS listener if it is not already enabled. For more information, see [How to configure WinRm for HTTPS](#).

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the certificate thumbprint will automatically be added to the registration information in the catalog.

SSL certificate thumbprint

The SSL certificate thumbprint is displayed here. Confirm that the certificate thumbprint matches the thumbprint of the certificate on the host that you are adding.

For Linux-based virtual machines:

Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

Use existing user

Enable to select a previously entered username and password for the provider.

Guest OS Username/Password

For some tasks (such as cataloging file metadata, file restore, and IP reconfiguration), credentials must be established for the associated VM. Enter the username and password, and ensure that the VM can be accessed from the IBM Spectrum Protect Plus appliance either through DNS or hostname.

The default security policy uses the Windows NTLM protocol, and the user identity follows the default *domain\name* format if the Hyper-V virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

6. To troubleshoot a connection to a hypervisor VM, use the **Test** function.

The **Test** function verifies the hypervisor tool settings and tests DNS settings between the IBM Spectrum Protect Plus appliance and the VM. Select a single VM, and then click **Select Options**. You must select **Catalog file metadata**. Select **Use existing user** to select a previously entered username and password for the resource. Alternately, enter a username in the **Guest OS Username** and password in the **Guest OS Password** fields if you have not previously entered the username and password for the resource. Click **Test**. For more information, see [“Testing the connection to a Hyper-V Server virtual machine” on page 244](#).

7. Click **Save**.

8. To configure additional options, click the **Policy Options** field that is associated with the job in the **SLA Policy Status** section. Set the additional policy options:

Pre-scripts and Post-scripts

Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and a script server where the script will run. Scripts and script servers are configured on the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Run inventory before backup

Run an inventory job and capture the latest data of the selected resources before starting the backup job.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backup operations for specific VMs or databases in the backup job definition. Separate multiple resources with a semicolon.

9. To save any additional options that you configured, click **Save**.

What to do next

After you define a backup job, complete the following action:

Action	How to
Create a Hyper-V restore job definition.	See “Restoring Hyper-V data” on page 251 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Starting jobs on demand” on page 433](#)

You can run any job on demand, even if the job is set to run on a schedule.

Excluding virtual disks from the SLA policy for a job

After you save a backup job definition, you can exclude individual virtual disks in a virtual machine from the SLA policy that is assigned to the job.

Before you begin

Excluding one or more virtual disks from a backup operation can impact the success of recovery. Consider the following information prior to excluding a virtual disk from a Hyper-V backup operation.

- For Instant Disk Restore, if a virtual disk is selected for a restore operation, an existing virtual machine (VM) is chosen as the destination. IBM Spectrum Protect Plus mounts the restored disk to the chosen destination VM.
- For Instant VM Restore, if the excluded virtual disk contains data that is necessary to boot the virtual machine, then the restored VM may fail to boot.
- For VMs with Windows-based guests, the restored VM may fail to boot if the disk on which the main operating system is installed, typically the C: drive, was excluded during the backup operation.
- For VMs with Linux-based guests, the restored VM may fail:
 - If a disk containing the boot or root (/) partition was excluded during backup.
 - If a disk containing a data (non-root) partition was excluded during backup, and the data volume did not have the 'nofail' option specified in /etc/fstab.

Procedure

To exclude virtual disks from the SLA policy that has been applied to a virtual machine:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Hyper-V**.
2. Select **VMs** in the **View** filter.
3. Click the link for the Hyper-V server that contains the virtual machine in the **Name** field.
4. Identify the virtual machine that contains the virtual disks to be excluded. If an SLA policy is already applied, it will be visible in the **SLA Policy** field.
5. Click the virtual machine name for the virtual machine to display the associated virtual disks.
6. Select the virtual disk that is to be excluded and then click **Select SLA Policy**.
7. Clear the selected SLA policy that is applied to the virtual disk in the SLA Policy pane.

Note: If multiple virtual disks are selected that are assigned to the same policy, the SLA policy will not appear as selected in the SLA Policy pane after clicking **Select SLA Policy**. Leave all SLA policies clear to exclude the virtual disks from the policy.

8. Click **Save**.

Restoring Hyper-V data

Hyper-V restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Before you begin

Complete the following tasks:

- Ensure that a Hyper-V backup job was run at least once. For instructions, see [“Backing up Hyper-V data”](#) on page 245.
- Ensure that the destination that you plan to use for the restore job is registered in IBM Spectrum Protect Plus. This requirement applies to restore jobs that restore data to original hosts or clusters.
- Ensure that the latest Hyper-V integration services are installed.

For Microsoft Windows environments, see [Supported Windows guest operating systems for Hyper-V on Windows Server](#).

For Linux environments, see [Supported Linux and FreeBSD virtual machines for Hyper-V on Windows](#).

- Ensure that the appropriate roles for restore operations are assigned to the affected users. Grant users access to hypervisors and backup and restore operations in the **Accounts** pane. Roles and associated permissions are assigned during user account creation. For instructions, see [Chapter 16, “Managing user access,”](#) on page 455 and [“Managing user accounts”](#) on page 466.
- Windows file indexing and file restore on volumes residing on dynamic disks is not supported.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- When restoring a virtual machine by using clone mode and by using the original IP configuration, ensure that credentials are established through the **Guest OS Username** and **Guest OS Password** options within the backup job definition.

About this task

If a Virtual Hard Disk (VHDX) is selected for a restore job, IBM Spectrum Protect Plus automatically presents options for an Instant Disk Restore job, which provides instant writable access to data and application restore points.

An IBM Spectrum Protect Plus snapshot is mapped to a target server where the snapshot can be accessed or copied as required. All other sources are restored by using Instant VM restore jobs, which can be run in the following modes:

Production mode

Production mode enables disaster recovery at the local site from primary storage or a remote disaster recovery site, replacing original machine images with recovery images. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data jobs that are associated with the virtual machine continue to run.

Test mode

Test mode creates temporary virtual machines for development, testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. Test machines are kept running while they are needed to complete testing and verification and are then cleaned up. Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines that are used for production. Virtual machines that are created in test mode are also given unique names and identifiers to avoid conflicts within your production environment.

Clone mode

Clone mode creates copies of virtual machines for use cases that require permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines that are created in clone mode are also given unique names and identifiers to avoid conflicts within your

production environment. With clone mode, you must be sensitive to resource consumption because clone mode creates permanent or long-term virtual machines.

Restriction: Moving from test mode to production mode is not supported for Hyper-V.

Procedure

To define a Hyper-V restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Hyper-V > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Hyper-V**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:

- a) Review the available sources, including virtual machines (VMs) and virtual disks (VDisks). You can expand a source by clicking its name.

You can also enter all or part of a name in the **Search for** box to locate VMs that match the search criteria. You can use the wildcard character (*) to represent all or part of a name. For example, vm2* represents all resources that begin with "vm2".

- b) Click the plus icon  next to the item that you want to add to the restore list next to the list of sources. You can add more than one item of the same type (VM or virtual disk).

To remove an item from the restore list, click the minus icon  next to the item.

- c) Click **Next**.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: <p>Backup Restores data that is backed up to a vSnap server.</p>

Option	Description
	<p>Replication Restores data that is replicated to a vSnap server.</p> <p>Object Storage Restores data that is copied to a cloud service or to a repository server.</p> <p>Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape).</p> <ul style="list-style-type: none"> Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resource restore or recurring restore

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Demo The demonstration site from which to restore snapshots. This menu item is available only if you updated the product from IBM Spectrum Protect Plus 10.1.6 or earlier.</p>

Option	Description
	<p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

- On the **Set destination** page, choose the instance to be restored for the selected source and click **Next**:

Original Host or Cluster

Select this option to restore data to the original host or cluster.

Alternate Host or Cluster

Select this option to restore data to a local destination that is different from the original host or cluster, then select the alternative location from the available resources.

In the **VM Folder Destination** field, enter the virtual machine folder path on the destination datastore. Note that the directory will be created if it does not exist. Use "/" as the root virtual machine folder of the targeted datastore.

- On the **Set datastore** page, take the following actions:
 - If you are restoring data to an alternate Hyper-V host or cluster, select the destination datastore and click **Next**.
 - If you are restoring data to the original Hyper-V host or cluster, this page is not displayed.
- On the **Set network** page, specify the network settings to use for each chosen source and click **Next**.
 - If you are restoring data to the original Hyper-V host or cluster, specify the following network settings:

Allow OS to define IP configuration

Select this option to allow your operating system to define the destination IP address. During a test mode restore operation, the destination virtual machine receives a new MAC address along with an associated NIC. Depending on your operating system, a new IP address can be assigned based on the original NIC of the virtual machine, or assigned through DHCP. During a production mode restore the MAC address does not change; therefore the IP address should be retained.

Use original IP configuration

Select this option to restore to the original host or cluster using your predefined IP address configuration. During the restore operation, the destination virtual machine receives a new MAC address, but the IP address is retained.

- If you are restoring data to an alternate Hyper-V host or cluster, complete the following step:
In the **Production** and **Test** fields, set virtual networks for production and test restore job runs. Destination network settings for production and test environments should point to different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks that are associated with test and production modes will be used when the restore job is run in the associated mode.
8. On the **Restore methods** page, select the restore method to be used for source selection. Set the Hyper-V restore job to run in **Production**, **Test**, or **Clone**. When you run in production mode, you have the option to select **Replace virtual disks only and retain virtual machine configuration**. Enabling this option replaces the storage in the virtual machine with the virtual disks from a previous virtual machine backup. This restore method maintains the virtual machine configuration replacing only the storage. When production restore is selected, the virtual machine to which the disk replace is applied must be powered off and only restoring to the original location is supported. Additionally, options such as overwriting the virtual machine and restoring based on tags are not available because the virtual machine is not being recreated. After the job is created, it can be run in production or clone mode through the **Job Sessions** pane.
 9. You can also change the name of the restored VM by entering the new VM name in the **Rename VM (optional)** field. The **Rename VM** option is not available when the production mode is used with the replace virtual disks option enabled. Click **Next** to continue.
 10. Optional: On the **Job Options (optional)** page, configure advanced options and click **Next**.

Make IA clone resource permanent

Enable this option to move the virtual disk to permanent storage and clean up temporary resources. This action is accomplished by starting a Live Migration operation for the resources in the background. The destination of the Live Migration operation is the VM Configuration Datastore. The Instant Access disk is still available for read/write operations during this operation.

Power on after recovery

Toggle the power state of a virtual machine after a recovery is run. Virtual machines are powered on in the order in which they are recovered, as set in the Source step. If **Use original IP configuration** is selected, the **Power on after recovery** option is not honored.

Restriction: Restored virtual machine templates cannot be powered on after recovery.

Overwrite virtual machine

Enable this option to allow the restore job to overwrite the selected virtual machine. By default, this option is disabled.

Continue with restore even if it fails

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the restore job stops if the recovery of a resource fails.

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow to overwrite and force cleanup of pending old sessions

Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Append suffix to virtual machine name

Enter a suffix to add to the names of restored virtual machines.

Prepend prefix to virtual machine name

Enter a prefix to add to the names of restored virtual machines. Click Save to save the policy options.

11. Optional: On the **Apply scripts** page, choose the following script options and click **Next**.

- Select **Pre-script** to select an uploaded script, and an application or script server where the prescript runs. To select an application server where the script will run, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Post-script** to select an uploaded script and an application or script server where the postscript runs. To select an application server where the script runs, clear the **Use Script Server** check box. Navigate to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns COMPLETED. If a postscript completes with a nonzero return code, the postscript task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a FAILED status.
12. Take one of the following actions on the **Schedule** page:
- To run an on-demand job, click **Next**.
 - To set up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
13. On the **Review** page, review your restore job settings and click **Submit** to create the job.
- On-demand jobs will begin immediately; recurring jobs will begin at the scheduled start time.

What to do next

After the job is complete, select one of the following options from the **Actions** menu on the **Jobs Sessions** or **Active Clones** sections on the **Restore** pane:

Cleanup

Destroys the virtual machine and cleans up all associated resources. Because this is a temporary virtual machine to be used for testing, all data is lost when the virtual machine is destroyed.

Clone (migrate)

Migrates the virtual machine to the datastore and virtual network that are defined as the test network.

Related tasks

[“Backing up Hyper-V data” on page 245](#)

Use a backup job to back up Hyper-V data with snapshots.

[“Adding a Hyper-V server” on page 243](#)

When a Hyper-V server is added to IBM Spectrum Protect Plus, an inventory of the server is captured, enabling you to complete backup and restore jobs, as well as run reports.

Backing up and restoring Amazon EC2 data

To protect Amazon EC2 data, first add an account for your EC2 instances in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for those instances.

To add an EC2 account to IBM Spectrum Protect Plus, access keys are required. Access keys are long-term credentials for an Identity and Access Management (IAM) user or the Amazon Web Services (AWS) account root user.

For information about how to create an IAM user with access keys and the permissions that are required for IBM Spectrum Protect Plus, see [“Creating an AWS IAM user” on page 257](#).

For increased security, it is recommended that the AWS account root user is not used for IBM Spectrum Protect Plus. For more information about the root user, refer to the [AWS Identity and Access Management User Guide](#).

EC2 data is stored in Amazon Web Services (AWS) Elastic Block Store (EBS) snapshots rather than the vSnap server. IBM Spectrum Protect Plus manages these snapshots for backup and restore operations.

Ensure that your EC2 environment meets the system requirements in “[Hypervisor \(Microsoft Hyper-V and VMware\) and cloud instance \(Amazon EC2\) backup and restore requirements](#)” on page 21.

Creating an AWS IAM user

To complete tasks in the IBM Spectrum Protect Plus user interface, IAM users must have access keys and required permissions.

About this task

You can use the AWS Management Console to create an IAM user by using the following steps. These steps are condensed from the steps that are documented in the [AWS Identity and Access Management User Guide](#) to show settings that are required for IBM Spectrum Protect Plus. For the complete and detailed steps for creating an IAM user, refer to this guide.

To create a user, you must have IAM administrative permissions.

Procedure

1. Sign in to the [AWS Management Console](#) and click **Services > IAM** to open the IAM Management Console.
2. In the console navigation panel, click **Users > Add user**.
3. Type the user name for the new user.
4. Select **Programmatic access** for the AWS access type.

This access type is required to create an access key, which is required by IBM Spectrum Protect Plus. IBM Spectrum Protect Plus does not require the access type **AWS Management Console access**.

5. Click **Next: Permissions**.
6. Click **Attach existing policies directly**, and then click **Create policy**.
The **Create policy** page opens in a new browser window.
7. Click the **JSON** tab and enter the following actions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ec2:DetachVolume",
        "ec2:AttachVolume",
        "ec2:DeregisterImage",
        "ec2>DeleteSnapshot",
        "ec2:DescribeInstances",
        "ec2:CreateVolume",
        "ec2:DescribeTags",
        "ec2:CreateTags",
        "ec2:RegisterImage",
        "ec2:DescribeRegions",
        "ec2:RunInstances",
        "ec2:DescribeSnapshots",
        "ec2:DescribeAvailabilityZones",
        "ec2:CreateSnapshots",
        "ec2:DescribeVolumes",
        "ec2:CreateSnapshot",
        "ec2:DescribeSubnets",
        "iam:PassRole"
      ],
      "Resource": "*"
    }
  ]
}
```

8. Click **Review Policy**.
9. Type a name and description (optional) for the policy that you are creating.
10. Review the **Summary** section to see the permissions that are granted by the policy.

11. Click **Create policy**.
12. Close the browser window and return to the window that contains the **Add user** page.
13. Select the policy that you created from the list of policies.
14. Optional: Set a permissions boundary.
15. Click **Next: Tags**.
16. Optional: Add metadata to the user by attaching tags as key-value pairs.
You can use tags to filter resources when you back up or restore EC2 data.
17. Click **Next: Review**.
18. Review your choices, and then click **Create user**.
A new window opens showing the user name, access key, and secret key.
19. To view the secret key, click show **Show** next to the secret key.
20. Click **Download.csv** to save the access key ID and secret access key to a CSV file on your computer.
Store the file in a secure location. You cannot access the secret access key again after this dialog box closes.
21. Click **Close** close the window.

What to do next

Add an account for EC2. To create an account, follow the instructions in [“Adding an Amazon EC2 account” on page 258](#).

Adding an Amazon EC2 account

When an Amazon EC2 account is added to IBM Spectrum Protect Plus, an inventory of the instances that are associated with the account is captured. You can then run backup and restore jobs and generate reports for the instances.

Before you begin

An access key is required to add an EC2 account. The access key enables IBM Spectrum Protect Plus to connect to and inventory EC2 instances for data protection. Access keys that are already entered in IBM Spectrum Protect Plus are provided in a selection list. If the access key that you want to use is not in the list, you must add the access key and security key. Ensure that you have the access key and secret key that you want to add.

Procedure

To add an EC2 account, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Amazon EC2**.
2. Click **Manage Accounts**.
3. Click **Add Account**.
4. Populate the fields in the **Account Properties** section:

Account Name

Enter a meaningful name to identify the access key that you select for the account.

Use existing access key

To specify a previously entered access key for the account, select this option and then select the key from the **Select a key** list.

If you do not select this option, complete the following fields to add a key.

Access Key

Enter the access key.

Secret Key

Enter the secret key.

5. Click **Save**.

IBM Spectrum Protect Plus confirms a network connection, adds the EC2 account to the database, and then catalogs the account instances.

If a message indicates that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connection.

What to do next

When you add an EC2 account to IBM Spectrum Protect Plus, an inventory is automatically run on each instance that is associated with the account. Instances must be detected to ensure that they can be backed up. You can run a manual inventory at any time to detect updates. For instructions about running a manual inventory, see [“Detecting Amazon EC2 instances” on page 259](#).

Related tasks

[“Backing up Amazon EC2 data” on page 259](#)

Use a backup job to back up data in an Amazon EC2 instance.

[“Restoring Amazon EC2 data” on page 261](#)

Use a restore job to restore EC2 data from a backup copy. For example, if data on an instance is lost or corrupted. You can define a job that restores data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations available.

Detecting Amazon EC2 instances

Amazon EC2 instances are automatically detected after an EC2 account is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the account was added.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Amazon EC2**.
2. In the list of EC2 accounts, select an account or accounts or click the link for an account to navigate to the regions or instances that you want to inventory.
The navigation is in the order account > region > instance.
3. Click **Run Inventory**.

Backing up Amazon EC2 data

Use a backup job to back up data in an Amazon EC2 instance.

Before you begin

Complete the following steps:

1. Ensure that the accounts to be backed up are added to IBM Spectrum Protect Plus. For more instructions, see [“Adding an Amazon EC2 account” on page 258](#).
2. Ensure that one or more SLA policies are configured for the EC2 instances. For more instructions, see [“Creating an SLA policy for Amazon EC2 instances” on page 205](#).
3. Ensure that IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the backup job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).
4. If an account is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To define an EC2 backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Amazon EC2**.
2. Select the instances to back up in the Amazon EC2 Backup pane by taking one of the following actions:
 - To select all instances that are associated with an EC2 account, select the check box for the account. Any instances that are added to this account are automatically assigned to the SLA policy that you choose.
 - To select instances by region or specific instances, click the account name and navigate to the region or instance. The navigation is in the order account > region > instance. If an instance does not have an assigned name, the instance ID is shown as the instance name.

To search for available instances, use the search function and toggle the displayed instances by using the **View** filter. Available options are **Instances** and **Tags**.

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup criteria to the job definition from the **SLA Policy Status** table.
4. Optional: To configure additional options for the SLA policies that you have added to the definition, in the **Policy Options** column of the **SLA Policy Status** table, click the clipboard icon  for an SLA policy and set the following options.

If the job is already configured, click the icon to edit the configuration.

Pre-scripts and Post-scripts

Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs. Windows-based machines support batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and a script server where the script will run. Scripts and script servers are configured by using the **System Configuration > Script** page.

To continue running the job if the script that is associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Run inventory before backup

Run an inventory job and capture the latest data of the selected instances before starting the backup job.

Exclude Resources

Exclude specific instances from the backup job by using single or multiple exclusion patterns. Resources can be excluded by using an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Exclude Resources by Tag

Exclude specific resources based on associated VM tags from the backup job. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns

support standard alphanumeric characters as well as the following special characters: - _ and *. Multiple filters may be separated with a semicolon.

Force Full Backup of Resources

This option is not used for EC2 backup operations.

5. Click **Save** to create the job definition.

The job will run as defined by the SLA policies that you selected. To run the job immediately, click **Jobs and Operations > Schedule**. Select the job and click **Actions > Start**.

Tip: When the job for the selected SLA policy runs, all instances that are associated with that SLA policy are included in the backup operation. To back up only selected instances, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single instance, select the instance and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
- To run an on-demand backup job for one or more instances, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job” on page 439](#).

What to do next

After you define an EC2 backup job, create an EC2 restore job definition.

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Restoring Amazon EC2 data” on page 261](#)

Use a restore job to restore EC2 data from a backup copy. For example, if data on an instance is lost or corrupted. You can define a job that restores data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations available.

[“Starting jobs on demand” on page 433](#)

You can run any job on demand, even if the job is set to run on a schedule.

Restoring Amazon EC2 data

Use a restore job to restore EC2 data from a backup copy. For example, if data on an instance is lost or corrupted. You can define a job that restores data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations available.

Before you begin

Complete the following tasks:

1. Ensure that an EC2 backup job was run at least once. For instructions, see [“Backing up Amazon EC2 data” on page 259](#).
2. Ensure that IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).

About this task

IBM Spectrum Protect Plus uses clone mode to create long-term copies of instances.

Procedure

To define an EC2 restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Virtualized Systems > Amazon EC2 > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Amazon EC2**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click an account in the list to show the instances that are available for restore operations. You can also use the search function to search for available instances. Enter all or part of a name to locate instances that match the search criteria. You can use the wildcard character (*) to represent all or part of a name.
Use the **View** filter to toggle displayed instances.
 - b) Click the plus icon  next to the instance that you want to use as the source of the restore operation.

You can select more than one instance from the list. However, all selected instances must be in the same region.

If the instance has attached volumes, you can navigate to the volumes and select them for the restore operation. You cannot select both instances and attached volumes.

The selected instances or attached volumes are added to the restore list next to the account list. To remove an item from the list, click the minus icon  next to the item.
 - c) Click **Next** to continue.
 3. Complete the fields on the **Source snapshot** page to select the instance snapshots that you want to restore and click **Next** to continue.
The fields that are shown depend on the number of instances that were selected on the **Select source** page.
 - If a single instance is selected, select the date range for the snapshots that you want to restore. The snapshots that are available for that date range are listed. Select the snapshot that you want to restore.
 - If multiple instances are selected, select the date range for the snapshots they you want to restore. The instances that have snapshots within that date range are listed. For each instance, select the restore point that you want to restore.
 4. On the **Set destination** page, specify the Availability Zone that you want to restore instances to and click **Next**:
 - Original Availability Zone**
Select this option to restore instances to the original Availability Zone.
 - Alternate Availability Zone**
Select this option to restore instances to an Availability Zone that is different from the original Availability Zone, and then select the alternate location from the available resources.

If you are restoring an attached volume, select the destination instance in the alternate Availability Zone and enter an optional device name in the **Destination Attachment** section.
 5. On the **Set network** page, change the subnet for each Availability Zone if you selected **Alternate Availability Zones** on the **Set destination** page. If you selected **Original Availability Zone**, no settings are provided on this page. Click **Next** to continue.
The Availability Zone subnet must be in the same region as the instances that are selected in step [“2” on page 262](#).

6. On the **Restore method** page, you can change the name of the restored instance by entering the new instance name in the **Rename Instance (optional)** field. Click **Next** to continue.
7. If you are running the restore job in advanced mode, you can set additional options as follows:

Power on after recovery

Toggle the power state of an instance after a recovery is run. Instances are powered on in the order in which they are recovered.

Continue with restore even if it fails

Toggle the recovery of an instance in a series if the previous instance recovery fails. If disabled, the restore job stops if the recovery of an instance fails.

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Restore instance tags

Enable this option to restore tags that are applied to instances through EC2.

Prepend prefix to instance name

Enter a prefix to add to the names of restored instances.

Append suffix to instance name

Enter a suffix to add to the names of restored instances.

8. Optional: On the **Apply scripts** page, choose the following script options and click **Next**.
 - Select **Pre-script** to select an uploaded script, and an application or script server where the prescript runs. To select an application server where the script will run, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Post-script** to select an uploaded script and an application or script server where the postscript runs. To select an application server where the script runs, clear the **Use Script Server** check box. Navigate to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns COMPLETED. If a postscript completes with a nonzero return code, the postscript task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a FAILED status.
9. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

The begins after you click **Submit**, and an **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  .

All running jobs are viewable in the **Jobs and Operations > Running Jobs** page.

Related tasks

[“Adding an Amazon EC2 account” on page 258](#)

When an Amazon EC2 account is added to IBM Spectrum Protect Plus, an inventory of the instances that are associated with the account is captured. You can then run backup and restore jobs and generate reports for the instances.

Restoring files

Recover files from snapshots that are created by IBM Spectrum Protect Plus backup jobs. Files can be restored to their original or an alternate location.

Before you begin

Note the following procedures and considerations before restoring a file:

- Review the file indexing and restore requirements in [“File indexing and restore requirements” on page 21](#).
- Run a backup job with catalog file metadata enabled. Follow these guidelines:
 - Ensure that credentials are established for the associated virtual machine as well as the alternate virtual machine destination through the Guest OS Username and Guest OS Password option within the backup job definition.
 - Ensure that the virtual machine can be accessed from the IBM Spectrum Protect Plus appliance either through DNS or hostname. In a Windows environment, the default security policy uses the Windows NTLM protocol, and the user identity follows the default *domain\name* format if the Hyper-V virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.
 - For a file restore to complete successfully, ensure that the user ID that is on the target machine has the necessary ownership permissions for the file that is being restored. If a file was created by a user that differs from the user ID that is restoring the file based on Windows security credentials, the file restore job fails.

About this task

Restrictions:

- Encrypted Windows file systems are not supported for file cataloging or file restore.
- File indexing and file restore are not supported from restore points that were copied to cloud resources or repository servers.
- When restoring files in a Resilient File System (ReFS) environment, restores from newer versions of Windows Server to earlier versions are not supported. For example, restoring a file from Windows Server 2016 to Windows Server 2012.
- File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a non-default local administrator is entered as the **Guest OS Username** when defining a backup job. A non-default local administrator is any user that has been created in the guest OS and has been granted the administrator role.

This occurs if the registry key `LocalAccountTokenFilterPolicy` in `[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` is set to 0 or not set. If the parameter is set to 0 or not set, a local non-default administrator cannot interact with WinRM, which is the protocol IBM Spectrum Protect Plus uses to install the Windows agent for file cataloging, send commands to this agent, and get results from it.

Set the `LocalAccountTokenFilterPolicy` registry key to 1 on the Windows guest that is being backed up with `Catalog File Metadata` enabled. If the key does not exist, navigate to `[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` and add a `DWord` Registry key named `LocalAccountTokenFilterPolicy` with a value of 1.

To help avoid issues that can result from time zone differences, use an NTP server to synchronize time zones across resources. For example, you can synchronize time zones for storage arrays, hypervisors, and application servers that are in your environment.

If the time zones are out of sync, you might experience errors during application registration, metadata cataloging, inventory, backup, or restore, or file restore jobs. For more information about identifying and resolving timer drift, see [Time in virtual machine drifts due to hardware timer drift](#)

Hyper-V considerations

Only volumes on SCSI disks are eligible for file cataloging and file restore.

Linux considerations

If data is located on LVM volumes, the *lvm2-lvmetad* service must be disabled because it can interfere with the ability of IBM Spectrum Protect Plus to mount and resign volume group snapshots or clones. To disable the service, complete the following steps:

1. Run the following commands:

```
systemctl stop lvm2-lvmetad
```

```
systemctl disable lvm2-lvmetad
```

2. Edit the `/etc/lvm/lvm.conf` and specify the following setting:

```
use_lvmetad = 0
```

If data resides on XFS file systems and the version of the `xfsp` package is between 3.2.0 and 4.1.9, the file restore can fail due to a known issue in `xfsp` that causes corruption of a clone or snapshot file system when its UUID is modified. To resolve this issue, update `xfsp` to version 4.2.0 or later. For more information, see [Debian Bug report logs](#).

Procedure

To restore a file, complete the following steps.

1. In the navigation panel, click **Manage Protection > Virtualized Systems > File Restore** .

2. Enter a search string to search for a file by name, and then click the search icon .

For more information about using the search function, see [Appendix A, “Search guidelines,” on page 481](#).

3. Optional: You can use filters to fine-tune your search across specific virtual machines, date range in which the file was protected, and virtual machine operating system types.

Searches can also be limited to a specific folder through the **Folder path** field. The **Folder path** field supports wildcards. Position wildcards at the beginning, middle, or end of a string. For example, enter `*Downloads` to search within the Downloads folder without entering the preceding path.

Note: Only file objects for which a snapshot was taken during the date range that is specified will be visible. For those objects, when the arrow is clicked beside the file object, all previous snapshots for that file object are displayed.

4. To restore the file by using default options, click **Restore**. The file is restored to its original location.
5. To edit options before restoring the file, click **Options**. Set the file restore options.

Overwrite existing files/folder

Replace the existing file or folder with the restored file or folder.

Destination

Select to replace the existing file or folder with the restored file or folder.

To restore the file to its original location, select **Restore files to original location**.

To restore to a local destination different from the original location, select **Restore files to alternative location**. Then select the alternate location from available resources by using the navigation menu or the search function.

Restriction: A file can be restored to an alternate location only if credentials were established for the alternate virtual machine through the **Guest OS Username/Password** option in the backup job definition.

Enter the virtual machine folder path on the alternate destination in the **Destination Folder** field. If the directory does not exist, it will be created.

Click **Save** to save the options.

6. To restore the file by using defined options, click **Restore**.

Related tasks

[“Backing up VMware data” on page 220](#)

Use a backup job to back up VMware resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data” on page 232](#)

VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Chapter 10. Protecting file systems

File systems that contain directories and files that you want to protect can be registered with IBM Spectrum Protect Plus. Select the file system servers and the drives that contain data that you want to protect. Microsoft Windows ReFS and NTFS file systems can be registered with IBM Spectrum Protect Plus so that you can set up backup jobs or regularly scheduled service level agreement (SLA) policies.

You can protect local file systems that are assigned to a drive letter. Clustered volumes and drive shares are not protected by IBM Spectrum Protect Plus.

Windows file systems

After you successfully register the machine that hosts the Microsoft Windows NTFS or ReFS file system with IBM Spectrum Protect Plus, you can start to protect your data on the listed volumes and drives. You can also create an on-demand backup of your file systems data, or set up service level agreement (SLA) policies to run regular scheduled backup jobs.

Ensure that your environment in which the file system is located, meets the minimum system requirements. For more information about the system requirements, see [“File system requirements” on page 22](#).

The IP address of the machine you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. Both must have a Windows Remote Management service that is listening on port 5985.

The fully qualified domain name must be resolvable and route-able from the IBM Spectrum Protect Plus appliance server and from the vSnap server.

Prerequisites for file systems

All prerequisites for using IBM Spectrum Protect Plus with file systems must be met before you start protecting your resources.

Requirements for working with file systems with IBM Spectrum Protect Plus are available here, [“File system requirements” on page 22](#).

Note: The user ID for registering Windows file servers can be set up with one of the following Windows configurations:

- The *Local System Administrator* user account with the User Account Control (UAC) security component set to Disabled. With this user you must open the Windows system **Control Panel > User Account Control Settings**, and move the slider to **Never notify**.
- A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled. With this user, you must open the Windows system **Local Security Policy**. From the **Security Settings** menu, choose **Local Policies > Security Options > User Account Control: Run all administrators in Admin Approval Mode policy**, and set this option to Disabled. Ensure that your Local Administrator Group includes the Log on as Service policy option.
- Do not use German Umlaut in the file names as it results in failure of inventory for the File Server.

Space prerequisites

Ensure that you have enough space on the machine that hosts the file system you are protecting. For more information about space requirements, see [“Space requirements for protecting file systems” on page 268](#). When you are restoring data to an alternative location, allow for extra space. No files are overwritten during the restore process. When files of identical names are found, both copies are retained.

Handling a security certificate for Windows

To secure access for protecting file system files with IBM Spectrum Protect Plus, you must create a certificate and manage its placement.

About this task

Note: If the restore service cannot load the certificate, files are deleted and a new self-signed certificate and key are created.

Tip: If the IBM Spectrum Protect Plus file systems agent has run, you will find a self-signed certificate and key in the following location: %LOCALAPPDATA%\FSPA\. If the agent has not run yet, follow the steps to create and move the self-signed certificate and key.

The administrator can access this directory at the following path:

C:\Users\Administrator\AppData\Local\

Procedure

1. Create a key and signed certificate for the client machine.
Neither the key or the certificate can have password protection as this affects the loading of files.
2. Create a directory folder called FSPA at a location like this %LOCALAPPDATA%\FSPA.
3. Copy the key and certificate and place them in the FSPA folder.
4. Copy the key and certificate in this folder.
5. Rename the key to localfspagent.key.
6. Rename the certificate to localfspagent.crt.

Space requirements for protecting file systems

Before you start backing up data that is stored on the registered file system, ensure that you have enough free disk space on the machine and in the vSnap repository for backup and restore operations.

Adding a file system

To start protecting the data on an ReFS or NTFS file system, you must add the host address where the file system is located. You can repeat the procedure to add every host that you want to protect with IBM Spectrum Protect Plus.

Before you begin

Restriction: In an IBM Spectrum Protect Plus environment, you can assign only one application server or file server per host. For example, if you register a host as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

Note: The user ID for registering Windows file servers can be set up with one of the following Windows configurations:

- The *Local System Administrator* user account with the User Account Control (UAC) security component set to Disabled. With this user you must open the Windows system **Control Panel > User Account Control Settings**, and move the slider to **Never notify**.
- A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled. With this user, you must open the Windows system **Local Security Policy**. From the **Security Settings** menu, choose **Local Policies > Security Options > User Account Control: Run all administrators in Admin Approval Mode policy**, and set this option to Disabled. Ensure that your Local Administrator Group includes the Log on as Service policy option.

About this task

To add a file system to IBM Spectrum Protect Plus, you must have the DNS name or the IP address of the machine, a user ID, and the password.

Procedure

1. In the navigation, expand **Manage Protection > File Systems**.
2. In the **File Systems** page, click **Manage file servers**, and click **Add file server** to add the host server.
3. In the **File server properties** section, enter the DNS name or the IP address of the machine in the **Host Address** field.
4. Obtain the certificate thumbprint and verify that the certificate thumbprint matches the thumbprint of the certificate on the host. Click **Get SSL certificate thumbprint**.

Get SSL certificate thumbprint

Get the SSL certificate thumbprint for the Windows-based host. You must complete this step when registering servers for the first time or if the certificate on the server changes. This setting will only be visible if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences”](#) on page 173.

The HTTPS listener must be enabled on the host. You must create a self-signed certificate and then enable the HTTPS listener if it is not already enabled. For more information, see [How to configure WinRm for HTTPS](#).

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the certificate thumbprint will automatically be added to the registration information in the catalog.

SSL certificate thumbprint

The SSL certificate thumbprint is displayed here. Confirm that the certificate thumbprint matches the thumbprint of the certificate on the host that you are adding.

5. Specify the type of user for the Windows server you are adding.
 - Use an existing user ID and password.
 - Enter a new user ID and password.

Note: The user ID for registering Windows file systems must be set up with one of the following Windows configurations:

- The Local System Administrator user account with the User Account Control (UAC) security component disabled. With this user, you must access the User Account Control Settings dialog in your Windows system **Control Panel**, and move the slider to **Never**.
- A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled. With this user you must access the Local Security Settings dialog on your Windows system and disable the **User Account Control: Run all administrators in Admin Approval Mode policy** setting. Ensure that your Local Administrator Group includes the **Log on as Service** policy option.

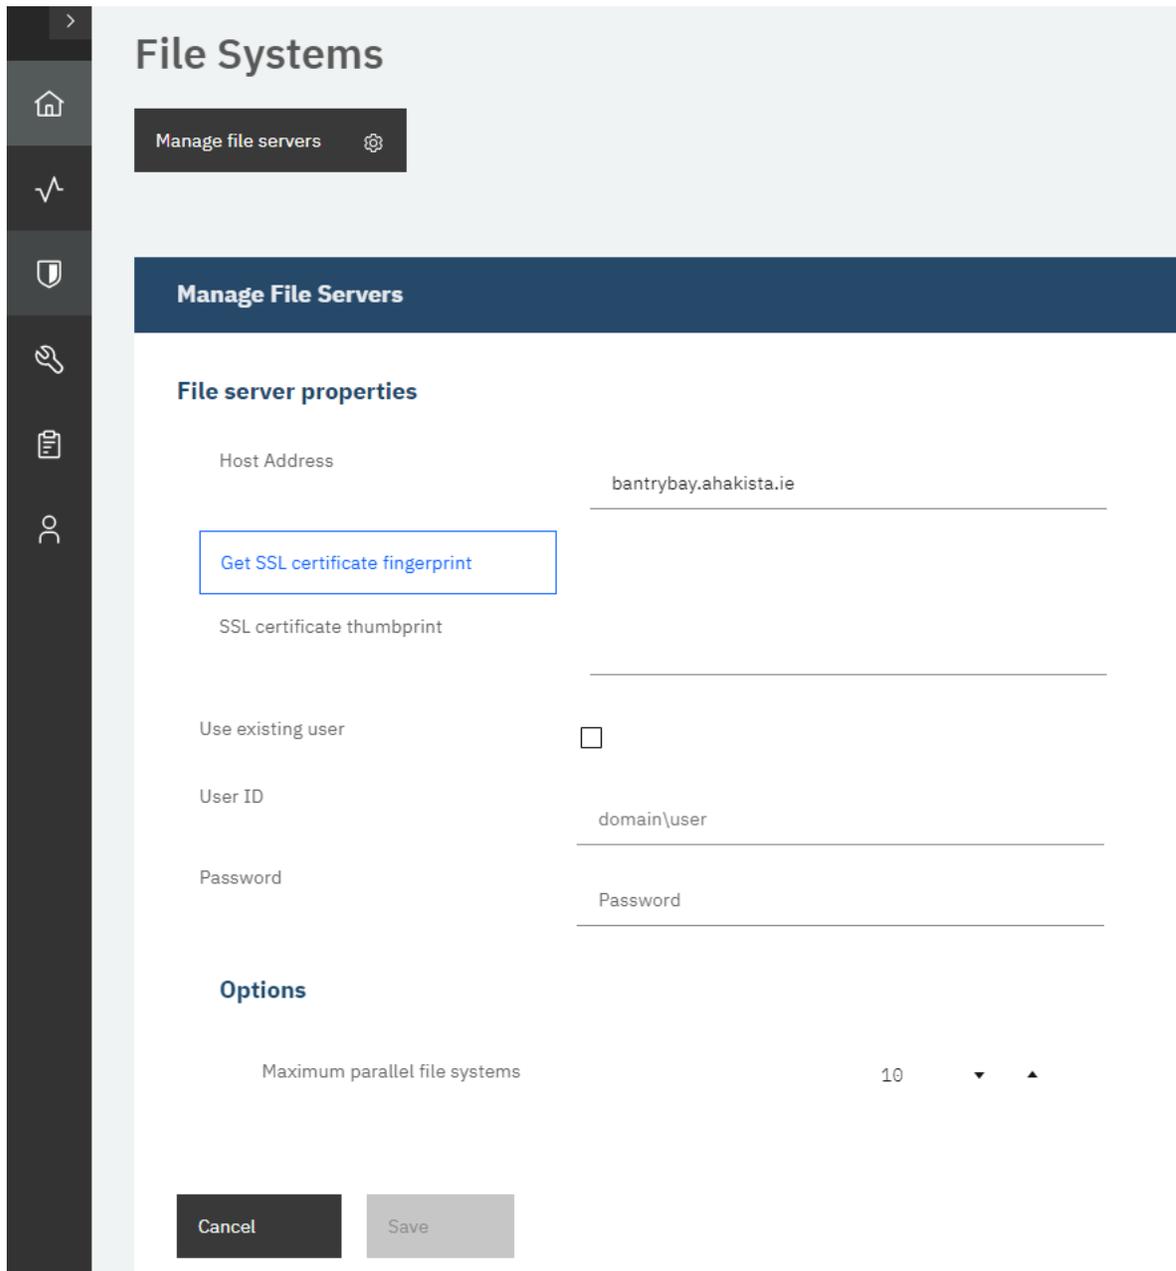


Figure 19. Managing agent users

Important: When you are entering the User ID, you do not need to enter the domain.

6. Set the maximum number of parallel file systems that are to be used for backing up data from the file system that is protected.

This setting applies to each file system on this host. Multiple resources can be backed up in parallel when the value of the option is set to more than 1. Multiple parallel file systems can speed up restore operations.

7. Save the form.

What to do next

After you add the file system host to IBM Spectrum Protect Plus, an inventory is automatically run to detect the relevant volumes and drives.

To verify that the drives and volumes are added, review the job log. Go to **Jobs and Operations**,



. Click the **Running Jobs** tab, and look for the Application Server Inventory log entry that corresponds to the inventory that was started.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

File systems must be detected to ensure that they can be protected. For instructions about running an inventory, see [Detecting file systems](#).

Running an inventory to detect file systems

After you add a file system to IBM Spectrum Protect Plus, an inventory to detect volumes, drives, and mount points is run automatically. The inventory detects, lists, and stores the file system resources that are found on the selected host, and makes the data available for protection with IBM Spectrum Protect Plus.

Before you begin

Ensure that you added the file system to IBM Spectrum Protect Plus. For instructions, see [Adding a file system](#).

Procedure

1. In the navigation panel, expand **Manage Protection > File Systems**.

Tip: To add file systems to the **Servers** pane, follow the instructions in [Adding a file system](#).

2. Click **Run Inventory**.

When the inventory is running, the text changes to show **Inventory In Progress**. You can run an inventory on any available file system server, but you can run only one inventory process at a time.

To view the job log, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the newest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name. If the job is not displayed, adjust the **Job History Period** to a longer time interval.

3. Click a server name to open a view that shows the volumes, drives, and mount points that are detected for that server. If any entries are missing from the **Servers** list, check your file systems and rerun the inventory. In some cases, certain entries are marked as ineligible for backup; hover over the entry to reveal the reason why.

Tip: To return to the list of servers, click the **Servers** hypertext.

Testing the file systems connection

After you add a file systems, you can test the connection. The test verifies communication with the server and the DNS settings between IBM Spectrum Protect Plus and the file systems server.

Procedure

1. In the navigation panel, click **Manage Protection > File Systems**.
2. In the **Microsoft Windows** window, click **Manage file servers**, and select the **Host Address** you want to test.

A list of the machine hosts that are available are shown.

3. Click **Actions** and choose **Test** to start the verification tests for physical network connection, remote access, and Windows privileges connections and settings.

The test report shows a list of the tests that were run. It consists of a test for the physical host network configuration, for the remote server installation on the host, and the Windows connections and privileges.

4. Click **OK** to close the test, and choose to rerun the test after you fix any failed tests.

Backing up file system data

Define regular backup jobs and specify options to run and create backup copies to protect your file system data.

Before you begin

During the initial backup, IBM Spectrum Protect Plus creates a new vSnap volume and NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus file system agent mounts the share on the server where the backup is to be completed.

Review the following procedures and considerations before you create a backup job definition:

- Add the file system servers that you want to back up. For the procedure, see [Adding a file system server](#).
- Configure a Service Level Agreement (SLA) Policy as described in this task.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,” on page 455](#).

A backup operation fails if the path is longer than 255 characters. If your paths are longer than 255 characters, you must enable longer paths by using the `Enable Win32 long paths` option in the Windows policy editor.

Note: IBM Spectrum Protect Plus does not protect file system shares, Microsoft cluster volumes or Microsoft cluster nodes.

About this task

The following steps describe how to back up resources that are assigned to an SLA policy. To run an on-demand backup job for one or more resources regardless of whether those resources are already associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job” on page 439](#).

Procedure

1. In the navigation panel, expand **Manage Protection > File Systems**.
2. Select a file system server to back up in the **Windows Backup** pane.
 - You can select an entire file system server by clicking the server name check-box. You can also select all file servers that are listed by clicking the check-box as shown. Any data added to the servers selected is automatically assigned to the SLA policy that you choose.
 - Or, you can select a specific drive or mount point from a specific file system server by clicking the server name, and choosing a drive or mount point from the list.
3. Click **Select Options** to specify files to be excluded from the backup job you are setting up. Alternatively, you can click **Modify Excluded Files** to leave the exclude rules as they are already defined. Click **Save** to commit your changes.

If you want to exclude all the files from a drive, you can specify the drive or a folder in a drive like this `Z:\test`. If you would like to exclude all files of a certain type from your backup job, you can specify that exclusion by using a string like this example `*.png`.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
 - If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
 - Options that are shown in yellow indicate that the option value has changed from the previously saved value.
 - To close the **Options** pane without saving changes, click **Select Options**.
4. Select the file systems server, drive, or mount point for backing up, and click **Select an SLA policy** to choose an SLA policy for that item.

You can choose from the following options: Gold, Silver, or Bronze. Each policy type has different frequencies and retention rates.

If you want to define a new SLA policy, select **Manage Protection > Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define your policy preferences. To edit an existing policy with custom retention and frequency rates, click the edit icon  and define your preferences. Click **Save** to commit your changes.

5. Click **Save** to save the SLA policy.

If you want to run the backup job immediately, click **Actions > Start**. The status in the log changes to show that the backup is Running.

What to do next

To view the status of your existing file system SLA policies, select **Manage Protection > Policy Overview** to view a summary of your protection.

Exclude rules syntax

You can define exclusion rules to exclude specified drives, directories, or files from backup jobs. After you define the rule, the specified items are not backed up as part of your SLA policy or as part of an ad hoc backup job. When you run a restore job, the drives, directories, or files that are specified in the exclude rules are not restored to the new copy.

Exclude rules can be defined for the entire Windows file systems application. Rules that define the excluded resources are inherited by each file system that is being protected. To define new rules for a particular file system instance, you can add a rule in the **File systems** window. Select the file system servers that you want to add the rules for. The new rules that you define for that file system backup job override the exclude rules that are set for Windows file systems.

For more information about defining a backup job, see [“Backing up file system data” on page 272](#).

To exclude a file, specify a rule as shown in the following example: `Z:\test\excludedFile.txt`.

To exclude all files in a folder, specify a rule as shown in the following example: `Z:\test*`.

To exclude a folder and its files, specify a rule as shown in the following example: `DIR Z:\excludedFolder`.

If you use global variables for certain directories, you can use those variables to exclude items regardless of their location. For example, if you wanted to exclude a directory that is called Blues owned by the registered user, you can specify `DIR %USERPROFILE%\Blues`. For more information about using global variables to exclude resources, see [Table 6 on page 275](#).

Table 5. Exclude rules syntax for Windows

Syntax	Syntax behavior
:\ :\	<ul style="list-style-type: none"> Indicates a file system and Windows drive. Must be included in all rules except for the FS rule. A rule cannot start or end with this syntax. A rule must start with a drive letter or wildcard followed by this sequence.
\	<ul style="list-style-type: none"> Indicates the next directory level. A rule cannot end with a backslash (\) character.
\...\	<ul style="list-style-type: none"> Indicates that the rule applies to all directories below this level. A rule cannot start or end with a \ . . . \ sequence. This sequence must be after the drive specification sequence.
*	<ul style="list-style-type: none"> This syntax is the wildcard for any character or number of characters. It is also used when no character is defined. A rule can start or end with this syntax. When used to indicate a drive letter, this syntax must represent one alphabetic character. This wildcard cannot be a backslash (\) character.
?	<ul style="list-style-type: none"> This syntax is used as a wildcard for any character for one occurrence only. A rule can start and end with this syntax. When this syntax is used to indicate a drive letter, it must be an alphabetic character between A and Z.
DIR	<ul style="list-style-type: none"> This syntax indicates a directory rule, but it does not exclude any files in the affected directory. This syntax must be a heading rule followed by a blank.
FS	<ul style="list-style-type: none"> Indicates that a full file system drive is excluded from the job. This syntax must be followed by a drive letter that can be a single character or a wildcard.
Spaces	<ul style="list-style-type: none"> Spaces are allowed in file names or directory names. A blank is not allowed before a backslash, \, or in a heading or trailing in a rule row. Spaces are validated as single characters.

Table 5. Exclude rules syntax for Windows (continued)

Syntax	Syntax behavior
Uppercase and lowercase text	Microsoft Windows is case-sensitive. Exclude rules ignore case.

Table 6. Exclude rules that use global variables

Syntax	Syntax behavior
DIR %PROGRAMDATA%	<ul style="list-style-type: none"> Indicates directory in the Windows ProgramData directory for the registered user. This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. For example, you can specify the following rule: DIR %PROGRAMDATA%\WinZip excludes the WinZip directory and all its content.
DIR %USERPROFILE%	<ul style="list-style-type: none"> Indicates directory in the Windows userProfile directory for the registered user. This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. For example, you can specify the following rule: DIR %USERPROFILE%\Elvis. This rule excludes the Elvis directory in that user's directory structure.
DIR %PROGRAMFILES%	<ul style="list-style-type: none"> Indicates a directory in the Windows Program Files directory for the registered user. This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. For example, you can specify the following rule: DIR %PROGRAMFILES%* to exclude all directories from the program files directory structure for the registered user.
DIR %WINDIR%	<ul style="list-style-type: none"> Indicates any specified Windows directory for the registered user. This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. An example is, DIR %WINDIR%\README. This rule excludes the README directory and all its content for the registered user.

Table 7. Valid exclude statements

Rule example	Result
:	This rule excludes all files from the file system root directory from all drives, but does not exclude the directories.
DIR *:*	This rule excludes all directories from all drives, but does not exclude the files in the root directory.

Table 7. Valid exclude statements (continued)

Rule example	Result
DIR E:\...*temp*	This rule excludes all directories that start with temp in the directory name in all directories of the E: drive.
DIR F:\Users\Bobby*	This rule excludes all content from the Bobby directory without excluding that directory itself. Files in the Bobby directory are excluded.
DIR F:\Users	This rule excludes all users who are listed in the Users directories and also excludes the Users directory.
DIR F:\Users\Bobby M?gee	This rule excludes all directories that match the name with a wildcard for one letter. This rule excludes users with names like Magee, Megee, Migege, and so on.
DIR F:\Users\Bobby Magee	This rule excludes the directory for the user who is defined, in this case Bobby Magee. With this rule, the directory for that user and all its content, including files and subfolders, are excluded.
F:\...*	This rule excludes all files from the F:\ drive, but it does not exclude the directories.
F:\Bobby.mp?	This rule excludes all files that match Bobby .mp? in the file system root directory, such as Bobby.MP3, Bobby.MP4, and so on.
F:\Bobby.txt	This rule excludes the file Bobby .txt in the file system root directory.
F:\Users\...*.mp3	This rule excludes all MP3 files for all users that are listed in the F drive.
F:\Users\Bobby\...*.mp3	This rule excludes all MP3 files from the user directory Bobby.
F:\Users\Bobby\...*music*\...*.mp?	This rule excludes all MP files in all directories that have the word music in the directory name for user Bobby. The excluded files are MP2, MP3, MP4, and so on.
F:\Users\John* DIR F:\Users\John*	This rule combination excludes all files and all subdirectories for the user John, but does not exclude the John directory itself.
F:\Users\John\tax\Tax_20???.pdf	This rule excludes all documents that match the pattern Tax_20 in the John\tax directory. Files that are named like these are excluded, TAX_2000.pdf, TAX_2019.pdf, and so on.
FS F	This rule excludes the file system F drive.
FS *	This rule excludes all drives in the file system.
FS ?	This rule excludes all drives.

Invalid exclude syntax

The following syntax is invalid for exclude rules:

- \no
- *
- *
- F:\no\
- DIR \no
- DIR F:\no\
- DIR *
- DIR F:*\

Tip:

To verify the results of an exclude rule, view the job log file. In the navigation panel, click **Jobs and Operations** and open the **Running Jobs** tab. In the **Application Server Backup** section, find the newest log entry.

Restoring file system data

To restore file system data from a vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. By using the File Systems File-Level Restore browser, you can select the file system resources to add to the job, and specify whether to restore data to another file system or to a different directory in the same file system on the same host.

Before you begin

Important: File are always restored to the same host machine.

- Ensure that at least one file system backup job was run successfully. For instructions about setting up a backup job, see [“Backing up file system data” on page 272](#).
- Ensure that appropriate IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).

Ensure that the system has sufficient space to allow the restore operation to complete. For more information about space requirements, see [Space requirements for file system protection](#). For more information about prerequisites and setup, see [Prerequisites for file system protection](#).

About this task



Attention: When you define the restore job, the **Run cleanup immediately on job failure** option is selected by default. This option must not be cleared unless you are instructed by the IBM Software Support team to do so.

Important: Restoring of files is always on the same host machine.

Procedure

1. In the navigation panel, expand **Manage Protection > File Systems** and click **Create job**.
2. Select **Restore**.
The **Restore** wizard opens.
3. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **file system** as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
4. On the **Select source** page, click a file system server to show the volumes that are available on that server. Select a volume by clicking the plus icon  next to that volume name . Only one file system volume from the backup can be added. Click **Next** to continue.
 5. On the **Source snapshot** page, select the snapshot that you want to restore. Click **Next** to continue. The available snapshots for the selected volume are listed with a timestamp, the associated SLA policy, and the source type that is available: backup, archive, or replication copy.
 6. On the **Review** page, review your selections for the restore job. If all selections are correct, click **Submit**, or click **Back** to edit the selections.
The **Active Resources** tab in Jobs and Operations is opened to show the active resource that is prepared when you exit the restore wizard.
Note: The active resource for the restore job that is submitted is not immediate and takes some time to display.
 7. Open the File Systems File-Level Restore browser by clicking **Open Browser** on the **Active Resources** tab.
 8. In the File Systems File-Level Restore browser, select the file system resources to add to the restore



job. Add items by clicking the add icon  next to the appropriate item.

9. You can specify more options for the restore operation as follows:
 - To overwrite an existing copy of the file or directory for the restore job, click **Overwrite**.
 - To specify an alternative location for the restore job, click **Options** and enter a valid Windows local volume path as the target.
 - To overwrite the existing copy of the file or directory at the alternative location, click **Overwrite**.
 - If you cannot find the file or directory that you want to add to the restore list, use the search capability. Search with options to specify searching for files or directories. Use wildcard symbols * to broaden the search for partial name strings.

Restriction: Network shares are not valid alternative locations for restore jobs.

10. Click **Restore** to start the restore process.
Existing files are overwritten only if the **Overwrite** option is selected. If files with identical names are detected during the operation, a timestamp is added to the new file and both files are stored at the restore location.
11. Optional: Monitor the progress of the restore operation in the **Restore Tasks** pane.
Tip: The restore process is not tracked on the IBM Spectrum Protect Plus **Jobs and Operations** page. Progress of the restore job is tracked in the File Systems File-Level Restore browser.

What to do next

When the restore job is completed, remove the active resource by taking the following actions:

1. In the navigation panel, click **Jobs and Operations > Active Resources**.
2. Select the active resource that you no longer require, and click **Cancel File System Restore**.

File Systems File-Level Restore browser

When you prepare a restore job for a specific file system, the active resource that is created can be viewed in the **File Systems File-Level Restore** browser so that you can define the items to be restored. Use the

browser to find and specify the directories or files that you want to restore from that file system. You can then specify an alternative location to direct the restored resources to a different location than the source.

Opening the File Systems File-Level Restore browser

After you click **Submit** in the Restore wizard, the restore job is prepared and the **Active Resources** tab in the **Jobs and Operations** page opens. To open the File Systems File-Level Restore browser, click the actions icon in the **Resources** table  and select **Open Browser** as shown.

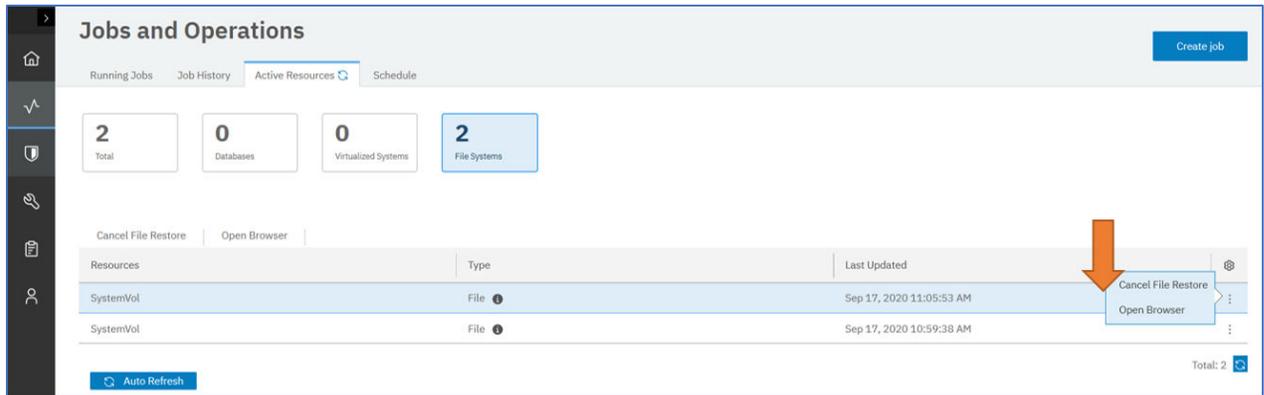
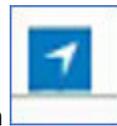


Figure 20. Opening the File Systems File-Level Restore from the Active Resources tab.

Adding resources to the restore job by using the File Systems File-Level Restore browser

To add specific file system resources to a restore job, navigate to the required file system, directories, or

files. Add items to the Restore List section by clicking the icon next to the file system item



Restoring file system resources to an alternative location

To clone or copy resources, and to restore those resources to a different location on the same file system, you can specify a valid Windows path as the target in the **Alternative Location** in the **Options** pane.

Monitoring a restore job

When you click **Restore** in the File Systems File-Level Restore browser, you can monitor the progress of the restore job in the **Restore Tasks** pane.

Chapter 11. Protecting data on cloud systems

Cloud systems such as Microsoft 365 is a cloud-based subscription service that can be registered with IBM Spectrum Protect Plus so that you can start to protect your data. Register Microsoft 365 with IBM Spectrum Protect Plus so that you can set up backup jobs or regularly scheduled service level agreement (SLA) policies to protect your data..

If you choose to protect Microsoft 365 with IBM Spectrum Protect Plus, you need to purchase IBM Spectrum Protect Plus for Microsoft 365 Entity ID Monthly License, Part Number D25ZELL. For more information about this entitlement, see [IBM Spectrum Protect Plus 10.1.5 announcement letter](#).

Microsoft 365

To protect Microsoft 365 email, calendars, contacts, and data on OneDrive cloud storage, you must register the Microsoft 365 application with Azure Active Directory.

Then, complete the following steps:

1. Deploy the Linux server that acts as the proxy host server.
2. Add a Microsoft 365 tenant and assign the cloud proxy server to the tenant.
3. Define a service level agreement (SLA) policy to create backup jobs.

If you choose to protect Microsoft 365 with IBM Spectrum Protect Plus, you need to purchase the IBM Spectrum Protect Plus for Microsoft Office 365 Entity ID Monthly License, Part Number D25ZELL. For more information about this entitlement, see the [IBM Spectrum Protect Plus 10.1.5 announcement letter](#).

Registering with Azure Active Directory

To protect a Microsoft 365 application, you must register the application with Azure Active Directory and grant appropriate permissions. When you register a new application with Azure Active Directory, the application credentials such as application ID and application secret are made available on the Azure Active Directory portal.

Before you begin

Take the following actions:

- Ensure that you have an active Microsoft 365 subscription.
- Ensure that you have a Microsoft 365 administrative user ID and password.

Tip: You can use an automated process to register and configure the Microsoft Azure Active Directory application. For instructions, see [technote 6437493](#).

Procedure

1. Go to the Microsoft 365 welcome page and sign in to your account by using your Microsoft 365 administrative user ID and password.
2. To open the Azure Active Directory admin center, in the left pane, click the ellipsis to expand the **Show all** menu, and then click **Admin centers > Azure Active Directory**.
3. To open your tenant dashboard, in the left pane of the Azure Active Directory admin center, click **Azure Active Directory**.
4. In the tenant dashboard menu, click **App registrations** and then click **New registration**.
5. To specify a user-facing name for the Microsoft 365 application, on the "Register an application" page, enter a name in the **Name** field.
6. Use the default options for the remaining fields, and click **Register**. The app registration is set up with the user-facing name that you entered.

7. To obtain the application (client) ID and directory (tenant) ID string, click **Azure Active Directory > tenant - App registrations > App name**. Then, copy the application ID string and directory ID. These strings will be required later, when you register the Microsoft 365 application with IBM Spectrum Protect Plus.
8. To create a client secret for this application ID, click **Certificates & secrets > New client secret**.
9. In the "Add a client secret" pane, enter any username in the **Description** field, and click **Add**. A client secret is generated, and the value is displayed in the "Client secrets" pane.
10. Copy the client secret to the clipboard by using the copy facility next to the **Client secret value** field. This character string is also used for registration with IBM Spectrum Protect Plus.
11. To add permissions for this application ID, click **API permissions > Add permissions**.
12. Specify permissions for each API in the following table by taking the following actions:
 - a) Select the **API name**, for example, Azure Active Directory Graph.
 - b) For the remaining permissions, select the **Application Permissions** type for each permission name for the APIs that are listed in the table.
 - c) For the permission name User.Read.All, select the **Delegated Permissions** type.
 - d) For the permission name full_access_as_app, select the **APIs my organization uses**, and enter Office 365 Exchange Online in the search field.

Remember: The Microsoft APIs view doesn't display the API name Office 365 Exchange Online by default.

API	Permission name
Microsoft Graph	Directory.Read.All
Microsoft Graph	User.Read.All
Office 365 Exchange Online	full_access_as_app
Microsoft Graph	Calendars.ReadWrite
Microsoft Graph	Contacts.ReadWrite
Microsoft Graph	Files.ReadWrite.All
Microsoft Graph	Mail.Read
Microsoft Graph	Mail.ReadWrite
Microsoft Graph	Mail.Send
Microsoft Graph	Sites.Read.All
Microsoft Graph	User.Read
Microsoft Graph	User.Read.All
Microsoft Graph	User.ReadWrite.All

13. To save the selected permissions, click **Grant admin consent for your organization name**, where *your organization name* specifies the name of your organization.

What to do next

Follow the instructions in [“Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus”](#) on page 282.

Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus

To ensure that the IBM Spectrum Protect Plus agent can connect to the Microsoft 365 tenant, you must register the Microsoft 365 tenant credentials, and the Linux server that acts as the proxy host server with

IBM Spectrum Protect Plus. This procedure is necessary to ensure that Microsoft 365 data can be backed up to IBM Spectrum Protect Plus.

Before you begin

- Ensure that you have a Linux system that can act as the cloud proxy server. IBM Spectrum Protect Plus deploys the backup agent on this server. For more information about the requirements, see [Microsoft 365 requirements](#).
- Ensure that the Microsoft 365 application is registered with Azure Active Directory. For instructions, see [“Registering with Azure Active Directory”](#) on page 281.

Procedure

1. In the navigation panel, expand **Manage Protection > Cloud Management > Microsoft 365**.

2. On the Microsoft 365 page, click **Manage application servers**, and then click **Add application server**.

3. On the Organization Properties page, complete the following fields:

- a. In the **Organization Name** field, enter the name of the organization that you set up in the Azure Active Directory admin center.

Note: Specify the Organization name in this format: *tenantname.onmicrosoft.com*, The names are not visible when you register the Azure application.

- b. In the **Tenant ID** field, enter the string from the **Directory (tenant) ID** field in the Azure Active Directory application registration.
- c. In the **Application ID** field, enter the string from the **Application (client) ID** field in the Azure Active Directory application registration.
- d. In the **Application Secret** field, enter the password string that was generated during the Azure Active Directory application registration.

4. On the Proxy Properties page, complete the following fields:

- a. In the **Host Address** field, enter the hostname or IP of the Linux server that is being used as the proxy host.
- b. Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

c. **Key type**

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

d. **Key fingerprint**

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

e. For host server authentication, select one of the following options:

- **User:** Select an existing user, or enter a user ID and the associated password.
- **SSH Key:** Select a Secure Shell (SSH) key from the drop-down list.

5. Click **Save**.

Results

When a proxy host is registered with IBM Spectrum Protect Plus, an inventory is run automatically on the Microsoft 365 organization, which returns the Microsoft 365 users in that resource.

Detailed process logs

The detailed process log is an additional Microsoft 365 process log file that can be useful for troubleshooting. This log is collected to track all backup and restore processes.

A detailed process log tracks the processes for each protected Microsoft 365 item. When you download the job log .zip file, you can view the detailed process log file along with standard diagnostic files.

Note: After you download the `joblog.zip` file, you can unzip the `diag.tar.gz` files to find the `Audit.log` file. This is the file with the Microsoft 365 processing information.

Detailed process log content and example

A detailed process log file includes the following information:

- Date and time of the operation.
- Operation type.
- Account that is associated with the operation.
- Indication of whether the event relates to OneDrive, a message, an calendar event, or a contact.
- Informational messages:
 - For OneDrive, the path and file name of the processed object is listed. If the operation is a redirected restore operation, that is indicated.
 - For messages, the date and time of the message is listed. If the operation is a redirected restore operation, any associated messages are listed.
 - For events, the subject of the event is listed.
 - For contacts, the name of the contact is listed.

Detailed process log example

The information in the detailed process log is provided in the following format:

```
[date time] [operation] [account] [relation] [message1] optional: [message2]
```

For example,

```
2020-02-13 19:15:27.805 Backup Completed username@example.com OneDrive
"my_new_document.pdf"
2020-02-13 19:13:46.754 Backup Completed username@example.com Message "1/20/2020 10:52:01
PM +01:00" "Welcome!"
2020-02-13 19:16:14.196 Backup Completed username@example.com Contact "John Smith"
2020-02-13 19:14:48.847 Backup Completed username@example.com Event "Monday meeting"
2020-02-13 19:18:22.544 Backup Failed username@example.com OneDrive
"my_folder\inventory.pdf"
2020-02-13 19:15:27.805 Restore Completed username@example.com OneDrive
"my_new_document.pdf" "my_new_document__2020-02-11_19_15.pdf"
2020-02-13 19:22:28.238 Backup Failed username@example.com OneDrive
"my_folder\inv\inventory.pdf"
```

Backing up Microsoft 365 data

After your Microsoft 365 organization is registered with IBM Spectrum Protect Plus, you can apply a service level agreement (SLA) policy to start protecting the Microsoft 365 data.

Procedure

1. In the IBM Spectrum Protect Plus navigation panel, expand **Manage Protection > Cloud Management > Microsoft 365**.
2. Select the checkbox for the organization.
3. Click **Select an SLA policy** and choose an SLA policy.
For more information about SLA policies, see [“Create backup policies” on page 98](#).
4. Save your choice. To define a new SLA or to edit an existing policy with custom retention periods or backup frequency rates, click **Manage Protection > Policy Overview**. In the "SLA policies" pane, click **Add SLA Policy**, and define policy preferences.

Tip: Some options in the **Policy Options** field in the **SLA Policy Status** section differ in availability based on backup type.

5. To run the policy outside the scheduled job, take the following actions:
 - a. To back up all organization data, select the checkbox for the organization.
 - b. To back up data from an account, click **Organization** and select the checkbox for the user name that is associated with the account.
 - c. To back up email, calendars, contacts, or OneDrive data for an account, click **Organization**, click area that you want to protect, and then click the user name and select the checkbox for the email, calendar, contacts, or OneDrive to back up.
6. Click **Run**. The status changes to **running** for the SLA and you can follow the progress of the job in the log.

Incremental forever backup for Microsoft 365

IBM Spectrum Protect Plus supports a backup strategy that is named *incremental forever*. Rather than scheduling periodic full backup jobs, this backup strategy requires only one initial full backup. Afterward, an ongoing sequence of incremental backup jobs occurs.

The incremental forever backup solution provides the following advantages:

- Reduces the amount of data that goes across the network
- Reduces data growth because all incremental backups contain only the objects that are new or changed since the previous backup job
- Reduces the duration of backup jobs

The IBM Spectrum Protect Plus incremental forever process includes the following steps:

1. The first backup job backs up all data from selected Microsoft 365 accounts.
2. All subsequent backup jobs back up only new or changed data from the selected accounts.

Restoring Microsoft 365 data

You can restore Microsoft 365 data from backup copies on vSnap servers or remote storage.

Before you begin

At least one Microsoft 365 backup job must have run successfully. For instructions about setting up a backup job, see [“Backing up Microsoft 365 data” on page 285](#).

About this task

The following restore modes are supported:

- Restore data to the original account
- Restore data to another account
- Restore data to a specified path

Procedure

1. In the navigation panel, expand **Manage Protection > Cloud Management > Microsoft 365**.
2. Click **Create job**.
3. Select **Restore**.
4. In the **Select source** pane, complete the following steps:
 - a) Click a source in the list to display the data that can be restored for the selected organization. You can also use the search function to search for available data and toggle the displayed data by using the **View** filter.
 - b) To select data to restore, click the Add to restore list icon  next to the data. You can select more than one item from the list. The selected items are added to the restore list. To remove an item from the source list, click the Remove from restore list icon  next to the data.
 - c) Click **Next** to continue.
5. On the "Source snapshot" page, select the restore type and the time when the data to be restored was backed up. Click **Next**.
6. On the "Select destination" page, complete the following fields, and click **Next** to continue.

Option	Description
Select a destination	Select the location to which data must be restored: Restore to original account Restores data to the original Microsoft 365 account Restore to another account Restores data to another Microsoft 365 account
Restore Path	Restores data to a selected directory path in the Microsoft 365 account

7. On the "**Job options**" page, if you want to run restore operations in parallel streams, specify a value in the **Max Parallel Streams** field. Click **Next**.
8. On the Review page, review your restore job settings.
9. To start the restore job, click **Submit**.

Results

A few moments after you click **Submit**, the on-demand restore job is added to the Running Jobs tab on the Jobs and Operations page. You can click the job record to display the details of the operation. You can also download the zipped log file by clicking **Download.zip**.

The account name for the restored data can be found in the log file for the restore operation. To locate the logs for a restore operation, in the navigation panel, click **Jobs and Operations** and then click the **Running Jobs** tab.

Chapter 12. Protecting databases

You must register the database applications that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the databases and resources that are associated with the applications.

Restriction: IBM Spectrum Protect Plus might create folders on application servers when applications are registered with IBM Spectrum Protect Plus. Folders created by IBM Spectrum Protect Plus must remain for the product to function properly. However, if you must remove a folder that was created by IBM Spectrum Protect Plus, unregister the application and IBM Spectrum Protect Plus will clean up the folders that are associated with the registration.

Do not assign more than one application per machine as an application server to a resource group. For example, if Microsoft SQL Server and Microsoft Exchange Server occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of the applications can be added as an application server to a given resource group.

Db2

After you successfully add your IBM Db2 instances to IBM Spectrum Protect Plus, you can start to protect your Db2 data. Create service level agreements (SLA) policies to back up and maintain Db2 data.

Tip: If your Db2 data is stored in a multi-partitioned environment with multiple hosts, you can protect your Db2 data across each host. Each host in the multi-partitioned environment must be added to IBM Spectrum Protect Plus so that all instances and databases are detected for protection. For more information, see [“Adding a Db2 application server”](#) on page 292.

The IP address must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. Both must have a Windows Remote Management service that is listening on port 5985.

The fully qualified domain name must be resolvable and routable from the IBM Spectrum Protect Plus appliance server and from the vSnap server.

Prerequisites for Db2

All prerequisites for the IBM Spectrum Protect Plus Db2 application server must be met before you start protecting Db2 resources with IBM Spectrum Protect Plus.

System requirements

Ensure that your IBM Spectrum Protect Plus Db2 application server environment meets the system requirements in [Db2 requirements](#).

Space prerequisites

Ensure that you have enough space on the Db2 database management system, in the volume groups for the backup operation, and on the target volumes for copying files during the restore operation. For more information about space requirements, see [Space requirements for Db2 protection](#). When you are restoring data to an alternative location, allocate extra dedicated volumes for the copy and restore processes. The data paths for table spaces and logs on the target host are the same as the paths on the original host. This setup is needed to allow copying of data from the mounted vSnap to the target host. Ensure that the protected Db2 database does not share the same local database directory with other Db2 databases.

Multi-partitioned Db2 environments

In order to protect Db2 multi-partitioned databases, the ACS backup mode must be set to parallel mode. To run parallel backup processing of partitions in your Db2 environment, ensure that one of the following prerequisites is met:

- The Db2 registry variable **DB2_PARALLEL_ACS** is set to YES, for example: **db2set DB2_PARALLEL_ACS=YES**.
- The Db2 registry variable **DB2_WORKLOAD** is set to SAP.

Restriction: The **DB2_PARALLEL_ACS** registry variable is available only in certain fix pack levels of Db2. If **DB2_PARALLEL_ACS** is not available in your version, you can choose to change **DB2_WORKLOAD** to SAP.

More configuration requirements

Ensure that your Db2 environment is configured to meet the following criteria:

- Db2 archive logging is activated, and Db2 is in recoverable mode.
- Ensure that the effective file size **ulimit -f** for the IBM Spectrum Protect Plus agent user and the Db2 instance user, is set to unlimited. Alternatively, set the value to a sufficiently high value to allow copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the Db2 instance to finalize the configuration.
- If you are running IBM Spectrum Protect Plus in an AIX or Linux environment, ensure that the installed sudo version is at the recommended level. For more information, see technote [2013790](#). Then, set sudo privileges as described in “Setting sudo privileges for Db2” on page 292.
- In a Linux environment, ensure that the Linux utility package **util-linux-ng** or **util-linux** package is current.
- Unicode characters in file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.
- The database table spaces, online logs, and the local database directory can be on one or separate dedicated logical volumes that are managed by either LVM2 or JFS2. For layout two examples, see the following pictures. In the first picture, two types of volume groups shown. In the second picture, all volumes for data and logs are on one volume group.

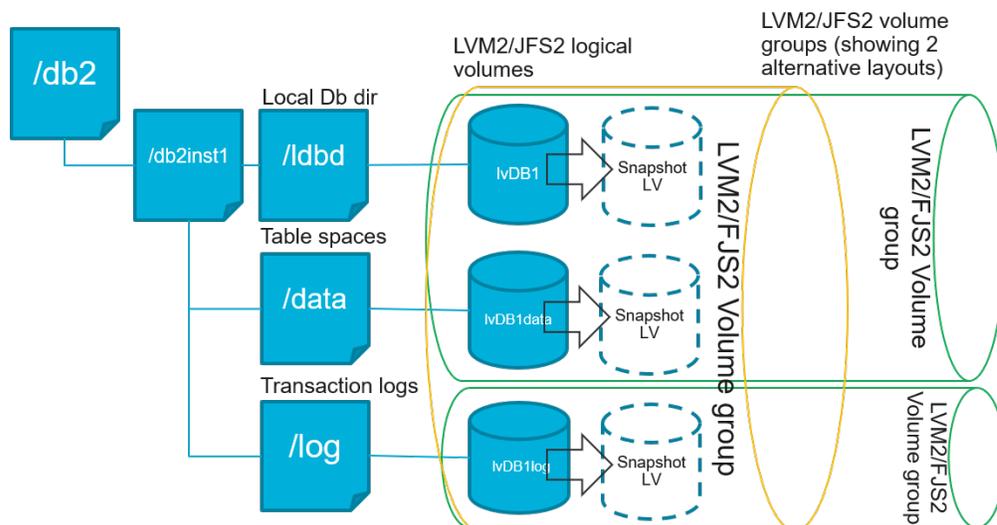


Figure 21. Logical volume layout examples

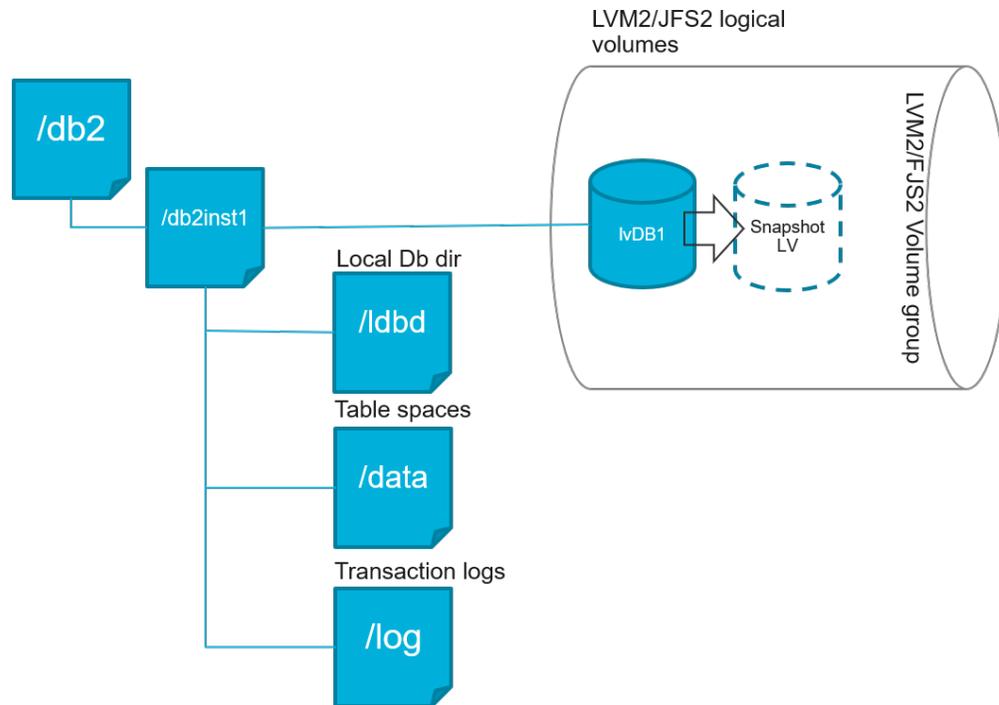


Figure 22. Single logical volume layout example

- Ensure that your Db2 logical volume setup does not include nested mount points.

Space requirements for Db2 protection

Before you start backing up Db2 databases, ensure you have enough free disk space on the target and source hosts, and in the vSnap repository. Extra free disk space is required on the volume groups on the source host for creating temporary Logical Volume Manager (LVM) snapshots of the logical volumes that the Db2 database and log files are stored on. To create LVM snapshots of a protected Db2 database, ensure that the volume groups with Db2 data have sufficient free space.

LVM snapshots

LVM snapshots are point-in-time copies of LVM logical volumes. They are space-efficient snapshots with the changed data updates from the source logical volume. LVM snapshots are created in the same volume group as the source logical volume. The IBM Spectrum Protect Plus Db2 agent uses LVM snapshots to create a temporary, consistent point-in-time copy of the Db2 database.

The IBM Spectrum Protect Plus Db2 agent creates an LVM snapshot which is then mounted, and is copied to the vSnap repository. The duration of the file copy operation depends on the size of the Db2 database. During file copying, the Db2 application remains fully online. After the file copy operation finishes, the LVM snapshots are removed by the IBM Spectrum Protect Plus Db2 agent in a cleanup operation.

For AIX, no more than 15 snapshots can exist for each JFS2 file system. Internal and external JFS2 snapshots cannot exist at the same time for the same file system. Ensure that no internal snapshots exist on the JFS2 volumes as these snapshots can cause issues when the IBM Spectrum Protect Plus Db2 agent is creating external snapshots.

For every LVM or JFS2 snapshot logical volume containing data, allow at least 10 percent of its size as free disk space in the volume group. If the volume group has enough free disk space, the IBM Spectrum Protect Plus Db2 agent reserves up to 25 percent of the source logical volume size for the snapshot logical volume.

LVM2 and JFS2

When you run a Db2 backup operation, Db2 requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system or a Journaled File System (JFS) for each logical volume with data or logs for the selected database. In Linux systems, the logical volumes are managed by LVM2 with `lvm2` commands. On AIX, the logical volumes are managed by JFS2 and created with the JFS2 snapshot command as external snapshots.

A software-based LVM2 or JFS2 snapshot is taken as a new logical volume on the same volume group. The snapshot volumes are temporarily mounted on the same machine that runs the Db2 instance so that they can be transferred to the vSnap repository.

On the Linux operating system, the LVM2 volume manager stores the snapshot of a logical volume within the same volume group. On the AIX operating system, the JFS2 volume manager stores the snapshot of a logical volume within the same volume group. For both, there must be enough space on the machine to store the logical volume. The logical volume grows in size as data changes on the source volume while the snapshot exists. In multi-partitioned environments, when multiple partitions share the same volume, an extra snapshot of the volume is created for each partition. Ensure that the volume group has sufficient free space for the required snapshots.

Setting sudo privileges for Db2

To use IBM Spectrum Protect Plus to protect your data, you must install the required version of the sudo program. For the Db2 application server, you must set up sudo in a specific way that might be different from other application servers.

Before you begin

To determine the correct version of sudo to be installed, see technote [2013790](#).

About this task

Set up a dedicated IBM Spectrum Protect Plus agent user with the required superuser privileges for sudo. This configuration enables the agent user to run commands without a password.

Procedure

1. Create an application server user by issuing the following command:

```
useradd -m <agent>
```

where `agent` specifies the name of the IBM Spectrum Protect Plus agent user.
2. Set a password for the new user by issuing the following command:

```
passwd <agent>
```
3. To enable superuser privileges for the agent user, set the `!requiretty` setting. At the end of the sudo configuration file, add the following lines:

```
Defaults:<agent> !requiretty
<agent> ALL=(ALL) NOPASSWD:ALL
```

If your sudoers file is configured to import configurations from another directory, for example `/etc/sudoers.d`, you can add the lines in the appropriate file in that directory.

Adding a Db2 application server

To start protecting your Db2 data, you must add the host address where your Db2 instances are located. You can repeat the procedure to add every host that you want to protect with IBM Spectrum Protect

Plus. If your Db2 environment is multi-partitioned with multiple hosts, you must add each host to IBM Spectrum Protect Plus.

About this task

To add a Db2 application server to IBM Spectrum Protect Plus, you must have the host address of the machine.

Procedure

1. In the navigation, expand **Manage Protection > Databases > Db2**.
2. In the **Db2** window, click **Manage application servers**, and click **Add application server** to add the host machine.
3. In the **Application Properties** section, enter the **Host Address**.
The host address is a resolvable IP address, or a resolvable path and machine name.
4. Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

5. Choose to specify a user or use an SSH key.
 - If you selected to specify a user, either select an existing user or enter a user ID and password.
 - If you are using an SSH key, choose the key from the menu.

Note: The user must have sudo privileges set up.

The screenshot shows the 'Manage Application Servers' configuration interface. At the top, there is a 'Manage application servers' button with a gear icon and a 'Create job' button. Below this is a dark blue header with the text 'Manage Application Servers'. The main section is titled 'Application Properties' and contains the following fields and controls:

- Host Address:** A text input field containing '77.00.999.12'.
- Get server key:** A blue button located below the Host Address field.
- Key type:** A text input field.
- Key fingerprint:** A text input field.
- User type:** Two radio buttons: 'User' (selected) and 'SSH Key'.
- Use existing user:** A checkbox that is currently unchecked.
- User ID:** A text input field containing 'domain\user'.
- Password:** A text input field with the placeholder text 'Password'.

At the bottom of the form, there are two buttons: 'Cancel' (dark grey) and 'Save' (light grey).

Figure 23. Managing agent users

Tip:

Db2 instances found are listed for each host. If your Db2 instance is partitioned, this information is listed with the host machine and the numbers of the partitions. For multi-host Database Partitioning Feature (DPF), the Db2 instance is displayed as a single unit.

6. Save the form, and repeat the steps to add other Db2 application servers to IBM Spectrum Protect Plus.

If your Db2 data is in a multi-partitioned environment with multiple hosts, you must add each host. Repeat the procedure for each Db2 host.

What to do next

After you add your Db2 application servers to IBM Spectrum Protect Plus, an inventory is automatically run on each application server to detect the relevant databases in those instances.

To verify that the databases are added, review the job log. Go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

Databases must be detected to ensure that they can be protected. For instructions about running an inventory, see [Detecting Db2 resources](#).

Detecting Db2 resources

After you add IBM Db2 application servers to IBM Spectrum Protect Plus, an inventory to detect all Db2 instances and databases is run automatically. The inventory detects, lists, and stores all the Db2 databases for the selected host, and makes the databases available for protection with IBM Spectrum Protect Plus.

Before you begin

Ensure that you added your Db2 application servers to IBM Spectrum Protect Plus. For instructions, see [Adding a Db2 application server](#).

Important:

Only short hostnames should be used in *db2nodes.cfg* for Db2. When hosts in the *db2nodes.cfg* are specified by their fully qualified names, which includes dots, cataloging of the database instance fails on inventory with the following error message:

```
Cataloging failed for server <servername>:  
org.springframework.data.mapping.model.MappingException:  
Map key <hostname> contains dots but no replacement was configured!  
Make sure map keys don't contain dots in the first place or configure an appropriate  
replacement!
```

About this task

Any Db2 partitions that are found in the inventory are listed for the Db2 instance. Partitions are listed by their partition number for each host appended to the host name in the **Instances** table.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Db2**.

Tip: To add more Db2 instances to the **Instances** pane, follow the instructions in [Adding a Db2 application server](#).

2. Click **Run Inventory**.

When the inventory is running, the button changes to show **Inventory In Progress**. You can run an inventory on any available application servers, but you can run only one inventory process at a time.

To view the job log, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

3. Click on an instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your Db2 application server and rerun the inventory. In some cases, certain databases are marked as ineligible for backup; hover over the database to reveal the reason why.

Tip: To return to the list of instances, click the **Instances** hypertext in the **Backup Db2** pane.

What to do next

To start protecting Db2 databases that are cataloged in the selected instance, apply a service level agreement (SLA) policy to the instance. For instructions about setting an SLA policy, see [Defining an SLA policy](#).

Testing the Db2 connection

After you add a Db2 application server, you can test the connection. The test verifies communication with the server and the DNS settings between IBM Spectrum Protect Plus and the Db2 server. It also checks for the correct sudo permissions for the user.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > Db2**.
2. In the **Db2** window, click **Manage Application Servers**, and select the **Host Address** you want to test.
A list of the Db2 application servers that are available are shown.
3. Click **Actions** and choose **Test** to start the verification tests for physical, remote and operating system connections and settings.
4. Click **OK** to close the test, and choose to rerun the test after you fix any failed tests.

Backing up Db2 data

Define regular Db2 backup jobs with options to run and create backup copies to protect your data. You can enable continuous backing up of archive logs so that you can restore a point-in-time copy with rollforward options if required.

Before you begin

During the initial backup, IBM Spectrum Protect Plus creates a new vSnap volume and NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus Db2 agent mounts the share on the Db2 server where the backup is to be completed.

Review the following procedures and considerations before you create a backup job definition:

- Add the application servers that you want to back up. For the procedure, see [Adding a Db2 application server](#).
- Configure a Service Level Agreement (SLA) Policy. For the procedure, see [Defining a Service Level Agreement backup job](#).
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,” on page 455](#).
- Inventory jobs should not be scheduled to run at the same time as backup jobs.
- Avoid configuring log backups for a single Db2 database with many backup jobs. If a single Db2 database is added to multiple job definitions with log backup enabled, a log backup from one job can truncate a log before it is backed up by the next job. This might cause point-in-time restore jobs to fail.

About this task

The following steps describe how to back up resources that are assigned to an SLA policy. To run an on-demand backup job for one or more resources regardless of whether those resources are already associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job” on page 439](#).

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Db2**.
2. Select a resource to back up.
 - Select an entire instance in the **Instances** pane by clicking the instance name check-box. Any databases added to this instance are automatically assigned to the SLA policy that you choose.

- Select a specific database in an instance by clicking the instance name, and choosing a database from the list of databases in that instance.
3. Click **Select Options** to enable or disable log backup, and to specify parallel streams to minimize time taken for large data movement in the backup operation. Click **Save** to commit the options.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

Select **Enable Log Backup** to back up archive logs, which allows point-in-time restore options and recovery options. For Db2 log backup settings information, see [Log backups](#).

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

When you save the options, those options are used for all backup jobs for this database or instance as selected.

4. Select the database or instance again, and click **Select SLA Policy** to choose an SLA policy for that database or instance.
5. Save the SLA options.

To define a new SLA or to edit an existing policy with custom retention and frequency rates, select **Manage Protection > Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define your policy preferences.

What to do next

When the SLA policy is saved, you choose to run an on-demand backup any time by clicking **Actions** for that policy, and selecting **Start**. The status in the log changes to show that the backup is Running.

Important: If the backup operation fails with an error, follow the procedure described in [“Troubleshooting failed backup operations for large Db2, MongoDB, and SAP HANA databases” on page 473](#).

Defining a service level agreement backup job

After your Db2 databases are listed for each of your Db2 instances, select and apply a service level agreement (SLA) policy to start protecting your data.

Procedure

1. From the navigation menu, expand **Manage Protection > Databases > Db2**.
2. Select a Db2 instance to back up all the data in that instance, or click the instance name to view the databases available for backing up. You can then select individual databases in the Db2 instance that you want to back up.

You can back up an entire instance with all of its associated data, or back up one or more databases.

3. Click **Select an SLA Policy** to select an SLA policy, and then click **Save**.

Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy

or edit an existing policy. For more information see [“Creating an SLA policy for databases and file systems”](#) on page 206.

4. Click **Select Options** to define options for your backup, such as enabling log backups for future recovery options, and specifying the parallel streams to reduce the time that is required to back up large databases. Save your changes.
5. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.

To read about more SLA configuration options, see [“Setting SLA configuration options for a backup job”](#) on page 298.

6. To run the policy outside of the scheduled job, select the instance or database. Click **Actions** and select **Start**.

The status changes to **Running** for your chosen SLA and you can follow the progress of the job in the job log.

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
- To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job”](#) on page 439.

To pause the schedule of an SLA, click **Actions** and choose **Pause Schedule**.

To cancel a job after it has started, click **Actions** > **Cancel**.

Setting SLA configuration options for a backup job

After you set up a service level agreement (SLA) for your backup job, you can choose to configure more options for that job. You can run scripts, exclude resources from the backup operation, and force a full base backup copy of a database if required.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job you are configuring, click the clipboard icon  to specify extra configuration options.
If the job is already configured, click on the icon to edit the configuration.
2. Click **Pre-Script** and define your pre-script configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.
3. Click **Post-Script** and define your post-script configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.

Scripts and script servers are configured on the **System Configuration** > **Script** page. For more information about working with scripts, see [Configuring scripts](#).

4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.

If this option is selected, the backup or restore operation is reattempted and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.

5. To exclude resources from a backup job, specify the resources to exclude from the job. Enter an exact resource name in the **Exclude Resources** field. If you are unsure of a name, use wildcard asterisks that are specified before the pattern (**text*) or after the pattern (*text**). Multiple wildcards can be entered with standard alphanumeric characters and the following special characters: - _ and *. Separate entries with a semicolon.
6. To create a full new backup of a resource, enter the name of that resource in the **Force full backup of resources** field. Separate multiple resources with a semicolon.

The full backup creates a full new backup of that resource and replaces the existing backup of that resource for one occurrence only. After the full backup completes, the resource is backed up incrementally as before.

Log backups

Archived logs for databases contain committed transaction data. This transaction data can be used to run a rollforward data recovery when you are running a restore operation. Using archive log backups enhances the recovery point objective for your data.

Ensure that you select the **Enable Log Backups** option to allow rollforward recovery when you set up a backup job or service level agreement (SLA) policy. When selected for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database. This backup creates a separate volume on the vSnap repository, which is mounted persistently on the Db2 application server. The backup process updates either **LOGARCHMETH1** or **LOGARCHMETH2** parameters to point to that volume for log archiving purposes. The volume is kept mounted on the Db2 application server unless the **Enable Log Backup** option is cleared and a new backup job is run.

Restriction: In Db2 multi-partitioned environments, the **LOGARCHMETH** parameters across partitions must match.

When either **LOGARCHMETH1** or **LOGARCHMETH2** parameters are set with a value other than OFF, you can use archived logs for rollforward recovery. You can cancel log backup jobs at any time by clearing the **Enable Log Backups** option: go to **Manage Protection > Databases > Db2**, select the instance and click **Select Options**. This change takes effect after the next successful backup job completes, and the **LOGARCHMETH** parameter value is changed back to its original setting.

Important: IBM Spectrum Protect Plus can enable log backup jobs only when the **LOGARCHMETH1** parameter is set to LOGRETAIN or if one of the **LOGARCHMETH** parameters is set to OFF.

If the **LOGARCHMETH1** parameter is set to LOGRETAIN.

IBM Spectrum Protect Plus changes the **LOGARCHMETH1** parameter value to enable log backups.

If either **LOGARCHMETH1** or **LOGARCHMETH2** parameters are set to OFF and the other is set to DISK, TSM, or VENDOR.

IBM Spectrum Protect Plus uses the **LOGARCHMETH** parameter that is set to off to enable log backups.

If both **LOGARCHMETH** parameters are set to DISK, TSM, or VENDOR.

This setting combination causes an error when IBM Spectrum Protect Plus attempts to enable log backups. To resolve the error, set one of the parameters to OFF, and run the backup job with the **Enable Log Backups** option selected.

Truncating archive log backups

IBM Spectrum Protect Plus automatically deletes older transactional logs after a successful database backup. This action ensures that the capacity of the log archive volume is not compromised by retention of older log files. These truncated log files are stored in the vSnap repository until the corresponding backup expires and is deleted. The retention of database backups is defined in the SLA policy that you select. For more information about SLA policies, see [“Defining a service level agreement backup job” on page 297](#).

IBM Spectrum Protect Plus does not manage the retention of other archived log locations.

For more information about Db2 settings, see [IBM Db2 documentation](#).

Restoring Db2 data

To restore Db2 data from the vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. You can choose to restore data to the original instance or to an alternative instance on a different machine, and specify recovery options, and save the job.

Before you begin

Important: For all restore operations, Db2 must be at the same version level on the source and target hosts. In addition to that requirement, you must ensure that an instance with the same name as the instance that is being restored exists on each host. This requirement applies when the target instance has the same name, and when the names are different. In order for the restore operation to succeed, both instances must be provisioned, one with original name and the other with the new name.

If your Db2 environment includes partitioned databases, the data of all partitions is backed up during regular backup jobs. All instances are listed in the backup pane. Multi-partitioned instances are shown with partition numbers and host names.

Before you create a restore job for Db2, ensure that the following requirements are met:

- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up Db2 data” on page 296](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- Restore jobs can create data in the IBM Db2 log directory. In some cases, if more than one restore job is run, data will remain in the log directory from the previous job. As a result, the next attempt to restore a database to the original location fails unless the log directory is purged.

For example, if the Db2 log directory is empty and a restore job runs with the options **Restore to original instance, Overwrite existing databases, and Recover until end of backup**, the restore job is successfully completed. If the job is followed by a second job with the options **Restore to original instance, Overwrite existing databases, and Recover until end of available logs**, this second restore attempt fails because the original restore job left data in the Db2 log directory.

Note: When you are restoring multi-partitioned databases to an alternative location, ensure that the target instance is configured with the same partition numbers as the original instance. All of those partitions must be on a single host. When you are restoring data to a new instance that is renamed, both instances required for the restore operation must be configured with the same number of partitions.

Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes table spaces, online logs, and the local database directory. Ensure that dedicated volumes with sufficient space are allocated to the file system structure. Db2 must be at the same version level on the source and target hosts for all restore operations, and an instance of the same name must exist on each host. For more information about space requirements, see [Space requirements for Db2 protection](#). For more information about prerequisites and setup, see [Prerequisites for Db2](#).

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Db2** and click **Create job > Restore**.

The Rrestore wizard opens.

- Optional: If you started the restore wizard from the **Jobs and Operations** page, click **Db2** as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
- On the **Select source** page, click a Db2 instance to show the databases in that instance. Choose a database by clicking the plus icon  for that database name. Click **Next** to continue.
 - In the **Source snapshot** page, choose the type of restore operation required.
 - On-Demand: Snapshot:** creates a once-off restore operation from a database snapshot. The job is not set to recur.
 - On-Demand: Point-in-Time:** creates a once-off restore operation from a point-in-time backup of the database. The job is not set to recur.
 - Recurring:** creates a recurring job that runs on a schedule and repeats.

Tip:

For an **On-Demand: Snapshot** you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job you can select to recover until the end of the available logs, or recover until a specific point-in-time.

- Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.

Option	Description
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.

Option	Description
	When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

6. Choose a **restore method** appropriate for the destination chosen for the restore operation. Click **Next** to continue.

- **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
- **Production:** In this mode, the Db2 application server first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. That copied data is then used to start the database.
- **Test:** In this mode, the agent creates a new database by using the data files directly from the vSnap repository.
- Add a database name when you are restoring the database to a different location and you want to rename the database.

Tip:

Production is the only **restore method** that is available for restore operations to the original location. Any options not appropriate for the restore operation that you selected are not selectable.

To restore data to the original instance, follow the instructions in [Restoring to the original instance](#). To restore data to an alternative instance, follow the instructions in [Restoring to an alternate instance](#).

7. Set the destination for the restore operation by choosing one of the following options. Click **Next** to continue.

- **Restore to original instance:** this option restores data to the original server and original instance.
- **Restore to alternate instance:** this option restores data to a different specified location, creating a copy of the data at that location.

If you are restoring data to an alternative location, choose an instance in the **Instance** table before you click **Next**. The alternative instance must be on a different machine; unsuitable instances are not available for selection. For multi-partition databases, the target instance must have the same set of partitions on a single machine.

8. In the **Job Options** page, select the recovery, application, and advanced options for the restore operation you are defining.

Tip:

Recovery options are not available for instant access restore jobs.

- **No Recovery.** This option skips any rollforward recovery after the restore operation. The database remains in a `Rollforward pending` state until you decide whether you want to run the rollforward operation manually.
- **Recover until end of backup.** This option recovers the selected database to its state at the time the backup was created. The recovery process uses the log files that are included in the Db2 database backup.
- **Recover until end of available logs.** This option is available only if the logs are backed up in the Db2 backup job definition. IBM Spectrum Protect Plus uses the latest restore point. A temporary restore point for log backups is created automatically so that the Db2 database can be rolled forward to the end of the logs. This recovery option is not available if you selected a specific restore point from the list. This option is available only when you are running an on-demand point-in-time restore job which uses the latest backup.

- **Recover until specific point-in-time.** This option includes all the backup data up to a specific point-in-time. This option is available only if you enabled log backups in your Db2 backup job definition. Configure a point-in-time recovery by a specific date and time, for example, Jan 1, 2019 12:18:00 AM. IBM Spectrum Protect Plus finds the restore points directly before and after the selected point-in-time. During the recovery process, the older data backup volume and the newer log backup volume are mounted. If the point-in-time is after the last backup, a temporary restore point is created. This recovery option is not available if you selected a specific restore point from the list. This option is available only when you are running an on-demand point-in-time restore job that uses the newest backup.

Tip: To skip optional steps in the restore wizard, select **Skip optional steps** and click **Next**.

9. Optional: In the **Job Options** page, select the application options for the restore operation you are defining.

Tip:

Application options are not available for instant access restore jobs.

- **Overwrite existing databases.** Choose this option to replace existing databases that have the same names during the restore recovery process. If this option is not selected, the restore job fails when databases with the same name are found during the restore operation. If you select this option, ensure that the Db2 log directory and the Db2 mirror log directory have no data.



Attention: Ensure that no other databases share the local database directory as the original database or that data is overwritten when this choice is selected.

- **Maximum Parallel Streams per Database.** You can choose to run the restore operation of data in parallel streams. This option is useful if you are restoring a large database.
 - **Specify the size of the Db2 database memory set in KB.** Specify the memory, in KB, to be allocated for the database restore on the target machine. This value is used to modify the shared memory size of the Db2 database on the target server. To use the same shared memory size at both the source server and the target server, set the value to zero.
10. Optional: In the **Job Options** page, select the advanced options for the restore operation you are defining.
 - **Run cleanup immediately on job failure.** This option enables the automatic cleanup of backup data as part of a restore if recovery fails. This option is selected by default. Do not clear this option unless instructed by IBM® Support for troubleshooting purposes.
 - **Continue with restores of other selected databases even if one fails.** This option continues the restore operation if one database in the instance fails to be restored successfully. The process continues for all other databases that are being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.
 - **Mount point prefix.** For instant access restore operations, specify the prefix for the path where the mount point is to be directed.
 11. Choose script options in the **Apply Scripts** page, and click **Next** to continue.
 - Select **Pre-Script** to select an uploaded script, and an application or script server where the pre-script runs. To select an application server where the script runs, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Post-Script** to select an uploaded script and an application or script server where the post-script runs. To select an application server where the script runs, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns COMPLETED. If a postscript completes with a nonzero return code, the postscript task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a FAILED status.

12. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.

13. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking

the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

To restore data to the original instance, follow the instructions in [Restoring to the original instance](#). To restore data to an alternative instance, follow the instructions in [Restoring to an alternate instance](#).

Restoring Db2 data to the original instance

You can restore a database backup to its original instance on the original host. You can restore to the latest backup or an earlier Db2 database backup version. When you restore a database to its original instance, you cannot rename it. This restore option runs a full production restoration of data, and existing data is overwritten at the target site if the **Overwrite existing databases** option is selected.

Before you begin

If your Db2 environment includes partitioned databases, the data of all partitions is backed up during regular backup jobs. All instances are listed in the backup pane. Multi-partitioned instances are shown with partition numbers and host names.

Before you create a restore job for Db2, ensure that the following requirements are met:

- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up Db2 data” on page 296](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- Restore jobs can create data in the IBM Db2 log directory. In some cases, if more than one restore job is run, data will remain in the log directory from the previous job. As a result, the next attempt to restore a database to the original location fails unless the log directory is purged.

For example, if the Db2 log directory is empty and a restore job runs with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of backup**, the restore job is successfully completed. If the job is followed by a second job with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of available logs**, this second restore attempt fails because the original restore job left data in the Db2 log directory.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Db2** and click **Create job > Restore**.
The Rrestore wizard opens.
2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **Db2** as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
3. On the **Select source** page, click a Db2 instance to show the databases in that instance. Choose a database by clicking the plus icon  for that database name. Click **Next** to continue.
 4. In the **Source snapshot** page, choose the type of restore operation required.
 - **On-Demand: Snapshot:** creates a once-off restore operation from a database snapshot. The job is not set to recur.
 - **On-Demand: Point-in-Time:** creates a once-off restore operation from a point-in-time backup of the database. The job is not set to recur.
 - **Recurring:** creates a recurring job that runs on a schedule and repeats.

Tip:

For an **On-Demand: Snapshot** you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job you can select to recover until the end of the available logs, or recover until a specific point-in-time.

5. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.

Option	Description
	When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore</p>

Option	Description
	operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

6. In the **Restore Method** page, choose **Production** for the restore operation.

In **Production** mode, the Db2 application server first copies the files from the vSnap repository volume to the target host. That copied data is then used to start the database.

Tip: Avoid entering a new database name when you are restoring a production operation to the original instance as it will not be implemented.

7. Set the destination for the restore operation to **Restore to original instance** to restore data to the original server. Click **Next** to continue.

8. Choose options as described in “Restoring Db2 data” on page 300.

9. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.

10. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking

the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

Restoring Db2 databases to an alternative instance

You can restore a Db2 database to another Db2 instance on an alternative host. You can also choose to restore a database to an instance with a different name and rename the database. This process creates an exact copy of the database on a different host in a different instance. If you are restoring a resource to an alternative location, you can restore the same resource multiple times without specifying different target hosts.

Before you begin

Important: For all restore operations, Db2 must be at the same version level on the source and target hosts. In addition to that requirement, you must ensure that an instance with the same name as the instance that is being restored exists on each host. This requirement applies when the target instance has the same name, and when the names are different. In order for the restore operation to succeed, both instances must be provisioned, one with original name and the other with the new name.

Before you create a restore job for Db2, ensure that the following requirements are met:

- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see “Backing up Db2 data” on page 296.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see Chapter 16, “Managing user access,” on page 455.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

- Restore jobs can create data in the IBM Db2 log directory. In some cases, if more than one restore job is run, data will remain in the log directory from the previous job. As a result, the next attempt to restore a database to the original location fails unless the log directory is purged.

For example, if the Db2 log directory is empty and a restore job runs with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of backup**, the restore job is successfully completed. If the job is followed by a second job with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of available logs**, this second restore attempt fails because the original restore job left data in the Db2 log directory.

Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes table spaces, online logs, and the local database directory. Ensure that dedicated volumes with sufficient space are allocated to the file system structure. Db2 must be at the same version level on the source and target hosts for all restore operations, and an instance of the same name must exist on each host. For more information about space requirements, see [Space requirements for Db2 protection](#). For more information about prerequisites and setup, see [Prerequisites for Db2](#).

Restriction: If data exists on the local database directory to which you are restoring the database backup to, and the **Overwrite existing databases** option is not selected, the restore operation fails. No other data can share the local database directory where the backup is restored. When the **Overwrite existing databases** option is selected, any existing data is removed and the local database directory on the alternative host.

Note: When you are restoring multi-partitioned databases to an alternative location, ensure that the target instance is configured with the same partition numbers as the original instance. All of those partitions must be on a single host. When you are restoring data to a new instance that is renamed, both instances required for the restore operation must be configured with the same number of partitions.

About this task

Ensure that the disk paths for the redirected restore operation include the instance name and the database name. The information is needed for all types of paths: database paths, container paths, storage paths, and log and mirror log paths.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Db2** and click **Create job > Restore**.

The Rrestore wizard opens.

2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **Db2** as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
3. On the **Select source** page, click a Db2 instance to show the databases in that instance. Choose a database by clicking the plus icon  for that database name. Click **Next** to continue.
 4. In the **Source snapshot** page, choose the type of restore operation required.
 - **On-Demand: Snapshot:** creates a once-off restore operation from a database snapshot. The job is not set to recur.
 - **On-Demand: Point-in-Time:** creates a once-off restore operation from a point-in-time backup of the database. The job is not set to recur.
 - **Recurring:** creates a recurring job that runs on a schedule and repeats.

Tip:

For an **On-Demand: Snapshot** you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job you can select to recover until the end of the available logs, or recover until a specific point-in-time.

5. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p>

Option	Description
	<p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

6. Choose a **restore method** appropriate for the destination chosen for the restore operation. Click **Next** to continue.

- **Production:** In this mode, the Db2 application server first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. That copied data is then used to start the database.
- **Test:** In this mode, the agent creates a new database by using the data files directly from the vSnap repository.
- **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
- Add a database name when you are restoring the database to a different location and you want to rename the database.

7. Set the destination for the restore operation to **Restore to alternate instance** to restore data to a different location, which you can select from the list of eligible locations. Click **Next** to continue.
When you are restoring to an alternative location, choose an instance in the **Instance** table before you click **Next**. Unsuitable target instances cannot be selected.
8. Choose options as described in [“Restoring Db2 data ” on page 300](#).
9. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.
If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.
10. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

Exchange Server

After you successfully register an Exchange application server, you can start to protect Microsoft Exchange data with IBM Spectrum Protect Plus. Define a service level agreement (SLA) policy to create backup jobs with specific schedules, retention policies, and scripts.

Prerequisites for Exchange Server

Ensure that all prerequisites for your Microsoft Exchange application are met before you start protecting Exchange databases with IBM Spectrum Protect Plus.

For more information, see [“Microsoft Exchange Server requirements” on page 22](#).

Virtualization support

IBM Spectrum Protect Plus supports Exchange Server running on a physical (bare metal) server, as well as in a virtualization environment. The following virtualization environments are supported:

- VMware ESX guest operating system
- Microsoft Windows Hyper-V guest operating system

Privileges

To help ensure that an Exchange agent can work in your IBM Spectrum Protect Plus environment, you must set up the appropriate privileges for the Exchange user account.

Role-based access control

You are required to register the Exchange Server with IBM Spectrum Protect Plus with an Exchange user who has local administrator privileges and the correct role-based access control (RBAC) permissions.

Also, for granular restore operations you are required to use an Exchange user who has local administrator privileges and the correct RBAC permissions.

To meet the minimum requirements for an Exchange user, complete the following steps:

1. Verify that the Exchange user is a member of a local Administrator group and has an active Exchange mailbox in the domain. Ensure that the Exchange server version level must be equal or higher to the Exchange version level of restoring mailbox.

By default, Windows adds the Exchange Organization Administrators group to other security groups, including the local Administrators group. For Exchange users who are not members of the Exchange Organization Management group, you must manually add the user account to the local Administrators group by taking one of the following actions:

- On the computer of the domain member, click **Administrative tools > Computer Management > Local Users and Groups tool**.
- On a domain controller computer that does not have a local Administrators group or Local Users and Groups tool, manually add the user account to the Administrators group in the domain: Click **Administrative tools > Active Directory Users and Computers tool**.

2. Set the role and scope.

- Verify that the Exchange user has the correct RBAC permissions.

You must assign the following management roles to each Exchange user who will complete mailbox restore operations:

- Active Directory Permissions
- ApplicationImpersonation
- Databases
- Disaster Recovery
- Mailbox Import Export
- Public Folders
- View-Only Configuration
- View-Only Recipients

Place users who complete mailbox restore tasks into an Exchange Server role group that contains these roles.

Exchange Server includes several built-in role groups. The Organization Management role group by default contains most, if not all, of the roles that are listed.

Place users who must complete multiple mailbox restore tasks into the Organization Management role group (ensuring that the group contains all of the listed roles).

Alternatively, you can place the user into another role group that you created or any other built-in role group that contains the roles that are listed. A user whose name is not in the Organization Management role group or subgroups might experience slower performance during restore operations.

Important: You can manage Exchange role groups by using the Exchange Admin Center (EAC) or Exchange Powershell Cmdlets *only* if your user name is authorized by the security policy in your organization.

- Management role scope

Ensure that the following Exchange objects are in the management role scope for the Exchange user:

- The Exchange Server that contains the required data
- The recovery database that is created by IBM Spectrum Protect Plus
- The database that contains the active mailbox
- The database that contains the active mailbox of the user who completes the restore operation

Encrypting File System

IBM Spectrum Protect Plus for Exchange requires that Encrypting File System (EFS) is enabled in the local or group domain policy, and a valid Domain Data Recovery Agent (DRA) certificate is available. If a custom group policy is defined and linked to the organizational unit, ensure that the Exchange server is part of the organizational unit.

Exchange certificates

Exchange digital certificates must be installed and configured for the mailbox browser to function during a granular restore operation. Ensure that the current Exchange certificates are installed and configured correctly in your environment.

Note: With Exchange 2016 and Exchange 2019, the Exchange Server is configured to use Transport Layer Security (TLS) by default. This TLS security encrypts communication between internal Exchange servers, and between Exchange services on the local server.

Adding an Exchange application server

When you register Exchange Server, an inventory of Exchange databases is added to IBM Spectrum Protect Plus. When the inventory is available, you can start to back up and restore your Exchange databases and run reports.

About this task

To register an Exchange application server, you need the IP address or host name.

Restriction: You can assign only one application server or file server per host. For example, if you register a host as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

Procedure

To add an Exchange application server, complete the following steps:

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. On the **Exchange** page, click **Manage Application Servers**, and then click **Add Application Server** to add the host system.
3. In the **Application Properties** form, enter the IP or host address in the **Host Address** field.
4. Obtain the certificate thumbprint and verify that the certificate thumbprint matches the thumbprint of the certificate on the host. Click **Get SSL certificate thumbprint**.

Get SSL certificate thumbprint

Get the SSL certificate thumbprint for the Windows-based host. You must complete this step when registering servers for the first time or if the certificate on the server changes. This setting will only be visible if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences”](#) on page 173.

The HTTPS listener must be enabled on the host. You must create a self-signed certificate and then enable the HTTPS listener if it is not already enabled. For more information, see [How to configure WinRm for HTTPS](#).

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the certificate thumbprint will automatically be added to the registration information in the catalog.

SSL certificate thumbprint

The SSL certificate thumbprint is displayed here. Confirm that the certificate thumbprint matches the thumbprint of the certificate on the host that you are adding.

5. Enter a user ID in the format of active directory domain and user account (domain\user), and the associated password.

This user must have the correct Exchange roles and privileges. For more information about Exchange privileges, see [“Privileges”](#) on page 312.

6. In the **Maximum concurrent databases** field, set the maximum number of databases per service level agreement (SLA) policy that can be backed up concurrently. The default is 10. Valid values are 1 - 99.

This value might be higher or lower than the number of databases that are associated with an SLA policy. For example, if an SLA policy has 10 associated databases and this value is set to 2, a backup operation occurs for only 2 of the 10 databases at the same time. As each backup operation completes, a second backup operation starts until all databases are backed up. If an SLA policy has 5 associated databases and this value is set to 10, all 5 database backup operations occur at the same time.

This option applies only to SLA policies that are associated with multiple databases. For SLA policies that are associated with only one database, this option provides no function.

The maximum number of concurrent database backup operations is limited by your environment. Some things to consider are the vSnap server configuration, network bandwidth, and the physical disk configuration of your IBM Spectrum Protect Plus server.

For guidance about tuning your IBM Spectrum Protect Plus environment for best performance, see the [IBM Spectrum Protect Plus Blueprints](#).

7. Click **Save**, and repeat the steps to add other Microsoft Exchange instances to IBM Spectrum Protect Plus.

Important: In a database availability group (DAG) environment, register all Exchange application servers in the DAG.

What to do next

When you add your Exchange application server to IBM Spectrum Protect Plus, an inventory is automatically run on each instance. Databases must be detected to ensure that they can be backed up, and you can run a manual inventory at any time to detect updates. For instructions about running a manual inventory, see [“Detecting Exchange databases by running an inventory”](#) on page 315. For instructions about setting up Exchange database backup jobs, see [“Defining a Service Level Agreement backup job”](#) on page 317.

Detecting Exchange databases by running an inventory

When you add your Exchange Server instances to IBM Spectrum Protect Plus, an inventory is run automatically. However, you can run an inventory on an Exchange application server manually at any time to detect updates and list all of the Exchange databases for each instance.

Before you begin

Ensure that you added your Exchange instances to IBM Spectrum Protect Plus. For instructions about adding an Exchange instance, see [“Adding an Exchange application server”](#) on page 314.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. Click **Run Inventory**.

When the inventory is running, the button label changes to **Inventory In Progress**. You can run an inventory on any available application server, but you can run only one inventory process at a time.
3. To monitor the inventory job, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.
4. When the inventory job is complete, on the **Exchange Backup** pane, click an Exchange instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your Exchange application server and rerun the inventory.

Tip: To return to the list of instances, click the **Instances** hypertext in the Exchange Backup pane.

Testing the Exchange connection

After you register a Microsoft Exchange application server and add it to the application server list, test the connection. The test verifies communication between IBM Spectrum Protect Plus and the host application server.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. On the **Exchange** page, click **Manage Application Servers**.
The Microsoft Exchange application servers that are available are shown.
3. Click **Actions** for the Microsoft Exchange application server that you want to test, and then click **Test**.
The test report shows you a list of the tests that ran and their status. Each test procedure includes a test of the physical host network configuration, a remote session test, and a test of Windows prerequisites such as user administrator privileges.
4. Click **OK** to close the test. Run the test again after you fix any issues.

Backing up Exchange databases

To protect Exchange databases, you can define a backup job that runs continuously to create incremental backups. You can also run on-demand backup jobs outside of the schedule.

Before you begin

Ensure that the application servers that contain the Exchange databases that you want to back up are registered with IBM Spectrum Protect Plus. For more information, see [“Adding an Exchange application server”](#) on page 314.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. On the **Exchange Backup** pane, click the Microsoft Exchange instance, and then select the database to back up.
Each database is listed by instance or database name, the applied SLA policy, and the eligibility for log backup.
3. Click **Run**.
The backup job begins, and you can view the details in **Jobs and Operations > Running Jobs**.
Tip: The **Run** button is only enabled for a single database backup, and the database must have an SLA policy applied.
To run an on-demand backup job for multiple databases that are associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job”](#) on page 439.
4. To run backup jobs for multiple databases, select the databases in the Exchange backup pane, and click **Select an SLA Policy**.
For more information about defining SLA policy backup jobs, and backup job options, see [“Defining a Service Level Agreement backup job”](#) on page 317.

Defining a Service Level Agreement backup job

When your Exchange databases are listed for each of your Exchange Server instances, select and apply a service level agreement (SLA) policy to start protecting your data.

About this task

IBM Spectrum Protect Plus supports single or multiple Exchange databases per Exchange backup job. Multiple database backup jobs run sequentially.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. Select an Exchange instance to back up all the data in that instance, or click an instance name, and then select individual databases that you want to back up.
3. Click **Select an SLA Policy** to select an SLA policy, and then click **Save**.
Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy or edit an existing policy. For more information see [“Creating an SLA policy for databases and file systems”](#) on page 206.
4. Click **Select Options** to define options for your backup, such as enabling log backups for future recovery options, and specifying the parallel streams to reduce the time that is taken to back up large databases. Save your changes.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
 - If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
 - Options that are shown in yellow indicate that the option value has changed from the previously saved value.
 - To close the **Options** pane without saving changes, click **Select Options**.
5. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.
For more information about SLA configuration options, see [“Setting SLA configuration options for a backup job”](#) on page 317.
 6. To run the policy outside of the scheduled job, select the instance or database and then click **Actions > Start**.
The status changes to **Running** for your chosen SLA. To pause the schedule, click **Actions > Pause Schedule**, and to cancel a job after it has started, click **Actions > Cancel**.

Setting SLA configuration options for a backup job

After you set up a service level agreement (SLA) for your backup job, you can choose to configure more options for that job. Extra SLA options include running scripts, excluding resources from the backup operation, and forcing a full base backup copy if required.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job that you are configuring, click the clipboard icon to specify additional configuration options.
2. To define a pre-script configuration, select **Pre-Script** and take one of the following actions:

- To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
 - To run a script on an application server, clear the **Use Script Server** check box, and choose an application server from the **Application Server** list.
3. To define a post-script configuration, select **Post-Script** and take one of the following actions:
- To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
 - To run a script on an application server, clear the **Use Script Server** check box, and choose an application server from the **Application Server** list.

Scripts and script servers are configured on the **System Configuration > Script** page. For more information about working with scripts, see [Configuring scripts](#).

4. Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails.

If this option is selected, the backup or restore operation is attempted and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not attempted and the script task status is reported as FAILED.

5. Specify resources to exclude them from the backup job. Enter an exact resource name in the **Exclude Resources** field. If you are unsure of a name, use wildcard asterisks that are specified before the pattern (**text*) or after the pattern (*text**). Multiple wildcards can be entered with standard alphanumeric characters and the following special characters: - _ and *. Separate entries with a semicolon.

6. If you want to create a full backup of a particular resource, enter the name of that resource in the **Force full backup of resources** field. Separate multiple resources with a semicolon.

A full backup replaces the existing backup of that resource for one occurrence only. After that, the resource is backed up incrementally as before.

7. Click **Save**.

Backing up Exchange database logs

You can back up the database transaction logs for Exchange databases. Exchange log backups are scheduled by using Windows Task Scheduler. When log backups are available, you can run a rollforward data recovery during a restore operation to ensure that the data is recovered to the latest possible point in time.

About this task

When log backups are enabled, a Task Scheduler task is created on the Exchange server. The task runs a backup operation of your Exchange log files according to the SLA policy.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. Click the Exchange Server instance that you want to protect, and then select the databases whose logs you want to back up.

Tip: The **Eligible for Log Backup** column shows the databases for which you can run log backups. If a database is registered as not eligible for log backup, a hover help explanation is provided.

3. Click **Select Options** and then select **Enable Log Backup**.

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

4. For **Repeats**, enter the frequency of the log backups in **Subhourly, Hourly, Daily, Weekly, Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.
5. Choose the **Start Time** and select the time for the log backups to begin, and then click **Save**.

Results

The database transaction logs are backed up to the vSnap server according to the selected frequency.

Restriction: The database logs are backed up on the preferred node only. Only one Exchange Server instance at a time can write log backups to the vSnap server.

Any log backup issues that occur are displayed in the Alert notifications in IBM Spectrum Protect Plus.

Backing up Exchange databases in a Database Availability Group

You can back up the mailbox databases in an Exchange Database Availability Group (DAG) and specify whether to use the active copy or a passive copy of the database for the backup. The Exchange servers in a DAG environment synchronize the data between active and passive copies for high availability. You can also use the incremental forever backup strategy to back up the mailbox databases in a DAG. For more information, see [“Incremental forever backup strategy” on page 320](#).

Before you begin

- Ensure that the system requirements are met. For more information about the requirements, see [Microsoft Exchange Server database backup and restore requirements](#).
- If you have a DAG cluster with Resilient File System (ReFS) and want to run an incremental backup, ensure that you created a dedicated DAG node with New Technology File System (NTFS). This node is responsible for maintaining passive copies of the mail databases. For information, see [Prerequisites of incremental backup](#).

Note:

If you use the incremental backup strategy, you can restore an Exchange Server DAG backup by using the following methods:

- Test Restore: You can restore the database to the original instance or to an alternative instance.
- Production Restore: You can restore a replicated database copy to an active database copy.
- Granular Restore: You can restore the database to the original instance or to an alternative instance.

Restriction: The granular restore method is not supported on Windows Server 2019 Core. For more information, see [Restrictions in Microsoft Exchange Server database backup and restore requirements](#).

About this task

By using the information from an inventory job, IBM Spectrum Protect Plus provides a DAG view that displays all of the databases in an Exchange DAG environment. Each database has an active copy on one server in the DAG, and one or more passive copies on the other servers. By default, scheduled backups are taken from the server that the database is active on, but you can select a different server to back up a passive copy of the database.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > Exchange**.
2. In the **Exchange Backup** pane, click the **View** menu and select **Database Availability Groups**.
3. Click the Exchange DAG that you want to view, and then select the databases to back up.
4. Click **Select Options**. In the **Backup preferred node** list, select the instance to run the backups on.

With the **Backup preferred node** option, you can select a passive copy of the database for the backup.

Important: To run an incremental forever backup, you must set the DAG node with NTFS that you created as a preferred node.

5. Click **Select an SLA Policy** and then select an SLA policy from the list.

6. To create the job definition by using default options, click **Save**.

The DAG databases are scheduled for backup jobs in accordance with the selected SLA policies and the preferred node choices.

7. To run the selected policy outside of the schedule, in the **SLA Policy Status** pane, click **Actions > Start**.

Incremental forever backup strategy

IBM Spectrum Protect Plus provides a backup strategy called *incremental forever*. Rather than scheduling periodic full backup jobs, this backup solution requires only one initial full backup. Afterward, an ongoing sequence of incremental backup jobs occurs.

Prior to 10.1.13, the IBM Spectrum Protect Plus Exchange database supports full database backup on both New Technology File System (NTFS) and Resilient File System (ReFS) volumes. However, the incremental forever backups were only supported for Exchange databases on NTFS volumes, and not for Exchange databases residing on ReFS volumes.

Beginning with 10.1.13, the incremental forever backups are supported for Exchange databases on ReFS volumes. When backing up a database, the initial backup operation creates a full database backup, and subsequent backups are incremental forever.

The incremental forever backup solution provides the following advantages:

- Reduces the amount of data that goes across the network
- Reduces data growth because all incremental backups contain only the blocks that changed since the previous backup
- Reduces the duration of backup jobs

The IBM Spectrum Protect Plus incremental forever process includes the following steps:

1. The first backup job creates a VSS snapshot of the Exchange application. As a result, the database files are in an application consistent state. The complete database files are copied to the vSnap location.
2. All subsequent backups create a VSS snapshot of the Exchange application. The database files are in an application consistent state. However, only the change blocks of the database files are copied to the vSnap location.
3. The backups are reconstructed at each point in time that a backup is performed, making it possible to recover the database from any single backup point.

Restoring Exchange databases

If data in an Exchange database is lost or corrupted, you can restore the data from a backup copy. Use the **Restore** wizard to set up a restore job schedule or an on-demand restore operation. You can define a job that restores data to the original instance or to an alternative instance, with different types of recovery options and configurations available.

Before you begin

Ensure that the following requirements are met:

- At least one Exchange backup job is defined and ran successfully. For instructions about defining a backup job, see [“Defining a Service Level Agreement backup job” on page 317](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is defining the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).

- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

Important: For granular restore operations, you must log on to the Exchange application server and use the Microsoft Management Console (MMC) GUI to complete mailbox batch restore and mailbox restore browser tasks.

Procedure

To restore data in an Exchange database, take one of the following actions:

- Restore a database to the original instance and location.
- Restore a database to the original instance with a different file location.
- Restore a database to an alternative instance.
- Restore mailbox data by using the granular restore function.
- Restore a database in a database availability group (DAG).

Restoring an Exchange database to the original instance

Restore an Exchange database to its original instance by using production mode or test mode. Choose between restoring the latest backup or an earlier Exchange database backup version.

Before you begin

Ensure that the following requirements are met:

- At least one Exchange backup job is defined and ran successfully.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is defining the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).

About this task

When you restore a database to its original location in production mode, you cannot rename it. This restore option runs a full production restore operation, and existing data is overwritten at the target site.

Procedure

To define an Exchange restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

- c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	<p>For on-demand restore operations, specify a range of dates to show the available snapshots within that range.</p>
Restore Point	<p>For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.</p>
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, choose from the following options:

- **Test.** In test mode, the agent creates a new recovery database by using the data files directly from the vSnap repository. This restore type might be used for testing purposes.
- **Production.** In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

For Test restore only, in the **New Database Name** field, enter the new name for the restored database. The **New Database Name** field is also displayed when you choose Production restore, but this is for restoring to a new database location on the original instance. For detailed instructions on this task, see [“Restoring an Exchange database to a new location on the original instance”](#) on page 325.

6. On the **Set destination** page, select **Restore to original instance** and click **Next**.

Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options

Choose from the following recovery options:

No Recovery

This option skips any rollforward recovery after the restore operation. The database remains in a Rollforward pending state until you decide whether you want to run the rollforward recovery manually.

Recover until end of backup

Restore the selected database to the state at the time the backup was created.

Recover until end of available logs

This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time

When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options

Set the application options:

Maximum Parallel Streams per Database

Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
 - If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.
The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.

Restoring an Exchange database to a new location on the original instance

You can restore an Exchange database to its original instance, but to a new location on the application server. Choose between restoring the latest backup or an earlier Exchange database backup version.

About this task

When you restore a database to its original instance by using a production restore operation, you can restore the database to a new file location on the application server with a new name for the restored database. In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates a new database by using the restored files.

Procedure

To define an Exchange restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for

Option	Description
	<p>the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
<p>Use alternate vSnap server for the restore job</p>	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
<p>Restore Location Type</p>	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane. Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.

Option	Description
	<p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. In the **Restore Method** page, click the **Production** restore option.

Tip: It is mandatory to select Production mode for this restore operation.

- a) In the **Name** field, expand the database name to see the path information for the existing database on the application server.
- b) In the **New Database Name** field, enter the new name for the restored database.
- c) In the **Destination Path** field, enter the new directory location for the database file on the server, including the .edb name, and the logs location.



Warning: The destination directories that you enter in the **Destination Path** field must already exist on the application host. If not, then create the necessary directories on the server before you complete the restore operation.

For example, for a database that is named Database_A, enter
C:\<new_destination_path>\Database_A.edb, and for the location of the logs, enter
C:\<new_logs_path>.

6. On the **Set destination** page, select **Restore to original instance** and click **Next**.

Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options

Choose from the following recovery options:

No Recovery

This option skips any rollforward recovery after the restore operation. The database remains in a Rollforward pending state until you decide whether you want to run the rollforward recovery manually.

Recover until end of backup

Restore the selected database to the state at the time the backup was created.

Recover until end of available logs

This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time

When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options

Set the application options:

Maximum Parallel Streams per Database

Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
 - If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.

Restoring an Exchange database to an alternative instance

You can select a Microsoft Exchange database backup and restore it to an Exchange Server instance on an alternative host. You can restore the database in production mode or test mode to the alternative instance.

Before you begin

Ensure that the following requirements are met:

- Enough disk space and allocated dedicated volumes are available for the copying of files.
- The file system structure on the source server is the same as the file system structure on the target server. This file system structure includes table spaces, online logs, and the local database directory.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: <ul style="list-style-type: none">• Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server.

Option	Description
	<p>Object Storage Restores data that is copied to a cloud service or to a repository server.</p> <p>Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape).</p> <ul style="list-style-type: none"> Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
<p>Use alternate vSnap server for the restore job</p>	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
<p>Restore Location Type</p>	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
<p>Select a location</p>	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p>

Option	Description
	If you are restoring data from a cloud or repository server, select a server from the Select a location menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, choose from the following options:

- **Test.** In test mode, the agent creates a new recovery database by using the data files directly from the vSnap repository. This restore type might be used for testing purposes.
- **Production.** In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.
 - a) In the **New Database Name** field, enter a new database name.
 - b) (Production restore only) Expand the database name to see the source and destination path information. In the **Destination Path** field, enter the directory location of the Exchange database file on the alternative host, including the .edb name, and the logs location.



Warning: The destination directories that you enter in the **Destination Path** field must already exist on the alternative host. If not, then create the necessary directories on the alternative host before you complete the restore operation.

For example, for a database that is named Database_A, enter
 C:\<new_destination_path>\Database_A.edb, and for the location of the logs , enter
 c:\<new_logs_path>.

6. On the **Set destination page**, choose **Restore to alternate instance**, select the target instance that you want to restore the database to and then click **Next**.
7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options

Choose from the following recovery options:

No Recovery

This option skips any rollforward recovery after the restore operation. The database remains in a Rollforward pending state until you decide whether you want to run the rollforward recovery manually.

Recover until end of backup

Restore the selected database to the state at the time the backup was created.

Recover until end of available logs

This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time

When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options

Set the application options:

Maximum Parallel Streams per Database

Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
 - If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.

Restoring individual mailbox items by using a granular restore operation

You can restore Exchange individual mailbox items by using a granular restore operation and the IBM Spectrum Protect Plus Microsoft Management Console (MMC) GUI.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. If RBAC permissions were not assigned, you might encounter configuration errors in the IBM Spectrum Protect Plus MMC GUI for each missing role.

Tip:

If you encounter role-based configuration errors in the IBM Spectrum Protect Plus MMC GUI, you can set the required permissions manually to resolve the errors (see [“Privileges” on page 312](#)), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step [“15” on page 336](#)).

About this task

To start a granular restore operation, complete preparatory steps in the IBM Spectrum Protect Plus GUI, and then log in to the Exchange application server. Then, use the IBM Spectrum Protect Plus MMC GUI to restore user mailbox data from the recovery database that is created by the granular restore operation. A granular restore operation can be used to complete the following tasks:

- You can restore selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode .pst file.

- You can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
- You can restore an archive mailbox or a part of the mailbox, for example, a specific folder.
- You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Source select** page, complete the following steps:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.

- b) Click the plus icon  next to the database that you want to use as the source of the restore operation.

Tip: You must select only one database for a granular restore operation. If you select multiple databases, the granular restore option will not be available on the **Restore method** page.

The selected source is added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

- c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:

Option	Description
	<ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane. Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.

Option	Description
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, click **Granular Restore**.
The recovery database name is displayed in the **New Database Name** field. The name consists of the existing database name with the suffix `_RDB`.
6. On the **Set destination** page, select **Restore to original instance** and click **Next**.
Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.
7. Optional: In the **Job Options** page, **Recover until end of backup** and **Run cleanup immediately on job failure** are selected by default. Click **Next** to continue.
Restriction: Do not clear the **Run cleanup immediately on job failure** option unless instructed by IBM Support for troubleshooting purposes.
8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
 - If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.
The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.
11. In the navigation panel, click **Jobs and Operations > Active Resources > Databases** to view the recovery database and mount point details.

Tip: Click the  icon to display an information message that describes the next steps for completing the granular restore task.

12. Connect to the Exchange application server instance by using Remote Desktop Connection (RDC) or Virtual Network Computing (VNC) if connecting remotely, or by logging on to the Exchange Server machine locally.

Tip: To log user and its mailbox with appropriate requirements, see [“Privileges”](#) on page 312.

The granular restore operation automatically installs and starts the IBM Spectrum Protect Plus MMC GUI on the application server. If the MMC GUI fails to start, start it manually by using the path that is provided in the **Active Resources** information message.

13. In the IBM Spectrum Protect Plus MMC GUI, click the **Protect and Recover Data** node, and select **Exchange Server**.
14. On the **Recover** tab for the Exchange Server instance, click **View > Mailbox Restore Browser** to view the mailbox from the recovery database.
15. Optional: Run the IBM Spectrum Protect Plus configuration wizard:
 - a) In the navigation panel, click **Dashboard > Manage > Configuration > Wizards > IBM Spectrum Protect Plus Configuration**.
 - b) In the **Actions** pane, click **Start**.
The configuration wizard runs the requirements check.
 - c) When the requirements checks have run, click the **Warnings** link next to **User Roles Check**.
 - d) On the message dialog box, to add any missing roles, click **Yes**.
 - e) On the configuration wizard, click **Next**, and then click **Finish**.
16. In the **Mailbox Restore Browser > Source** tree, click the mailbox that contains the items you want to restore, which enables you to browse the individual folders and messages.

Choose from the following actions to select the folder or message to restore.

<i>Table 8. Previewing and filtering mailbox items</i>	
Task	Action
Preview mailbox items	<ol style="list-style-type: none"> a. Select a mailbox item, such as Inbox, to display its contents in the preview pane. b. Click an individual item in the preview pane, such as an email message, to view the message text and details. c. If an item contains an attachment, click the attachment icon to preview its contents.

Task	Action
Filter mailbox items	<p>Use the filter options to narrow the list of folders and messages to restore:</p> <ol style="list-style-type: none"> Click Show Filter Options and Add Row. Click the down arrow in the Column Name field and select an item to filter. You can filter by folder name, subject text, and other options. <p>Restriction: You can filter public mailbox folders only by the Folder Name column.</p> <p>When you select All Content, the mailbox items are filtered by attachment name, sender, subject, and message body.</p> In the Operator field, select an operator: Contains. In the Value field, specify a filter value. To specify additional filtering criteria, click Add Row. Click Apply Filter to filter the messages and folders.

17. When you have selected the mailbox item to restore, in the **Actions** pane, click the restore task that you want to run. Choose from the following options:

- **Restore Folder to Original Mailbox**
- **Restore Messages to Original Mailbox**
- **Save Mail Message Content**

Tip: If you click **Save Mail Message Content**, a Windows Save File window is displayed. Specify the location and message name and click **Save**.

When you choose the restore option, the **Restore Progress** window opens and shows the progress of the restore operation, and the mailbox item is restored.

18. To restore a mailbox item to another mailbox or .pst file, complete the following steps.

Note: You can also restore a complete mailbox to another mailbox or .pst file.

Choose from the actions in the following table:

Task	Action
Restore a mailbox item (or a mailbox) to a different mailbox	<ol style="list-style-type: none"> On the Actions pane, click Open Exchange Mailbox. Enter the alias of the mailbox to identify it as the restore destination. Drag the source mailbox item (or mailbox) to the destination mailbox on the results pane. <p>Restriction: You cannot drag mail items or subfolders in the Recoverable Items folder to a destination mailbox.</p>

Table 9. Restoring a mailbox item to another mailbox or .pst file (continued)	
Task	Action
Restore a mailbox item (or mailbox) to an Outlook personal folders (.pst) file	<p>a. On the Actions pane, click Open non-Unicode PST File.</p> <p>b. When the Open File window opens, select an existing .pst file or create a .pst file.</p> <p>c. Drag the source mailbox item (or mailbox) to the destination .pst file on the results pane.</p> <p>Restriction: You can use the Mailbox Restore Browser view only with non-Unicode .pst files.</p>
Restore a Public Folder	<p>Select this action to restore a public folder to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore.</p> <ul style="list-style-type: none"> To restore a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). <p>You can also restore all or part of a public folder to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p>

19. In the **Actions** pane, click **Close Exchange Mailbox** or **Close PST File** to close the destination mailbox or .pst file.

Tip: You can enable the Microsoft Management Console to gather diagnostic information to assist in problem determination related to restore operations. The process gathers configuration files, trace files, and overall diagnostics of the MMC GUI. For more information, see the following technote: [Enabling diagnostic information in the IBM Spectrum Protect Plus MMC GUI](http://www.ibm.com/support/docview.wss?uid=ibm10882270)(<http://www.ibm.com/support/docview.wss?uid=ibm10882270>).

20. When the restore operation for the individual items is finished, return to IBM Spectrum Protect Plus and click **Jobs and Operations > Active Resources > Databases**.
21. Select the resource, and then click **Cancel File Restore** to end the granular restore process.

Restoring mailboxes by using a granular restore operation

You can restore Exchange mailboxes by using a granular restore operation and the IBM Spectrum Protect Plus Microsoft Management Console (MMC) GUI.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. If RBAC permissions were not assigned, you might encounter configuration errors in the IBM Spectrum Protect Plus MMC GUI for each missing role.

Tip:

If you encounter role-based configuration errors in the IBM Spectrum Protect Plus MMC GUI, you can set the required permissions manually to resolve the errors (see “Privileges” on page 312), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step “15” on page 342).

About this task

To start a granular restore operation, complete preparatory steps in the IBM Spectrum Protect Plus GUI, and then log in to the Exchange application server. Then use the IBM Spectrum Protect Plus MMC GUI to restore user mailbox data from the recovery database that is created by the granular restore operation. A granular restore operation can be used to complete the following tasks:

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode .pst file.
- You can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
- You can restore an archive mailbox or a part of the mailbox, for example, a specific folder.
- You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Source select** page, complete the following steps:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation.

Tip: You must select only one database for a granular restore operation. If you select multiple databases, the granular restore option will not be available on the **Restore method** page.

The selected source is added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.

Option	Description
	<p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	<p>For on-demand restore operations, specify a range of dates to show the available snapshots within that range.</p>
Restore Point	<p>For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.</p>
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, click **Granular Restore**.
The recovery database name is displayed in the **New Database Name** field. The name consists of the existing database name with the suffix **_RDB**.
6. On the **Set destination** page, select **Restore to original instance** and click **Next**.
Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.
7. Optional: In the **Job Options** page, **Recover until end of backup** and **Run cleanup immediately on job failure** are selected by default. Click **Next** to continue.
Restriction: Do not clear the **Run cleanup immediately on job failure** option unless instructed by IBM Support for troubleshooting purposes.
8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:

- If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.
The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.
 11. In the navigation panel, click **Jobs and Operations > Active Resources > Databases** to view the recovery database and mount point details.

Tip: Click the  icon to display an information message that describes the next steps for completing the granular restore task.

12. Connect to the Exchange application server instance by using Remote Desktop Connection (RDC) or Virtual Network Computing (VNC) if connecting remotely, or by logging on to the Exchange Server machine locally.

Tip: To log user and its mailbox with appropriate requirements, see [“Privileges” on page 312](#).

The granular restore operation automatically installs and starts the IBM Spectrum Protect Plus MMC GUI on the application server. If the MMC GUI fails to start, start it manually by using the path that is provided in the **Active Resources** information message.

13. In the IBM Spectrum Protect Plus MMC GUI, click the **Protect and Recover Data** node, and select **Exchange Server**.
14. On the **Recover** tab for the Exchange Server instance, select **View > Mailbox Restore**.
A list of user mailboxes from all databases that are included in the backup is displayed.
15. Optional: Run the IBM Spectrum Protect Plus configuration wizard:
 - a) In the navigation panel, click **Dashboard > Manage > Configuration > Wizards > IBM Spectrum Protect Plus Configuration**.
 - b) In the **Actions** pane, click **Start**.
The configuration wizard runs the requirements check.
 - c) When the requirements checks have run, click the **Warnings** link next to **User Roles Check**.
 - d) On the message dialog box, to add any missing roles, click **Yes**.
 - e) On the configuration wizard, click **Next**, and then click **Finish**.
16. Select one or more mailboxes from the recovery database to restore. Mailboxes are listed by Mailbox Name, Alias, Server, Database, and Mailbox Type.

You can restore only user mailboxes that are located in the recovery database.

Tip: Mailboxes from other databases are shown in this view for informational purposes only. If the mailbox that you want to restore is not in the recovery database, use this view to determine which Exchange database the user mailbox was assigned to. You can then run the granular restore task again for that database.

17. To complete the restore operation, in the **Actions** pane, click one of the following restore options.

<i>Table 10. Restore options</i>	
Option	Action
Restore Mail to Original Location	Restore mail items to their location at the time of the backup operation.

<i>Table 10. Restore options (continued)</i>	
Option	Action
Restore Mail to Alternate Location	<p>Restore the mail items to a different mailbox.</p> <ul style="list-style-type: none"> On the Alternate Mailbox Options window, enter the Mailbox alias name. <p>Tip: If deleted mail items or tasks are flagged in the Recoverable Items folder of a mailbox, the items are restored with the flag attribute to the Flagged Items and Tasks view in the target mailbox.</p>
<p>Restore Mail to non-Unicode PST file</p> <p>Restriction:</p> <ul style="list-style-type: none"> This option is available only for Exchange Server 2013. Each folder can contain a maximum of 16,383 mail items. 	<p>Restore mail items to a non-Unicode personal folders (.pst) file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location. Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory.</p> <p>If the .pst file exists, the file is used. Otherwise, the file is created.</p>
Restore Mail to Unicode PST file	<p>Restore mail items to a Unicode .pst file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location.</p> <p>Tip:</p> <p>You can enter a standard path name (for example, c:\PST\mailbox.pst) or a UNC path (for example, \\server\c\$\PST\mailbox.pst). When you enter a standard path, the path is converted to a UNC path. If the UNC is a non-default UNC path, enter the UNC path directly.</p> <p>Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory. If the .pst file exists, the file is used. Otherwise, the file is created.</p>

Table 10. Restore options (continued)	
Option	Action
Restore Public Folder Mailbox	<p>Restore a public folder mailbox to an online public folder mailbox.</p> <p>In the Folder to be restored field, enter the name of the public folder that you want to restore:</p> <ul style="list-style-type: none"> • To restore a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. • To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox.</p>
Restore Mail to Archive Mailbox	<p>This action applies to a primary mailbox or an archive mailbox. Select this action to restore all or part of either type of mailbox to the original archive mailbox or to an alternative archive mailbox.</p> <p>You can filter the archive mailbox and restore a specific mailbox folder. In the Folder to be restored field, enter the name of the folder in the archive mailbox that you want to restore.</p> <ul style="list-style-type: none"> • To restore a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. • To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). <p>In the Target archive mailbox field, specify the archive mailbox destination.</p>
Exclude recoverable mail items while restoring the mailbox	<p>Apply this action if you are restoring an online, public folder, or archive mailbox to an original mailbox, alternative mailbox, or to a Unicode .pst file.</p> <p>Specify a value of Yes to exclude the mail items in the Recoverable Items folder in mailbox restore operations. No is the default value.</p>

Tip: You can enable the Microsoft Management Console to gather diagnostic information to assist in problem determination related to restore operations. The process gathers configuration files, trace files, and overall diagnostics of the MMC GUI. For more information, see the following technote: [Enabling diagnostic information in the IBM Spectrum Protect Plus MMC GUI](http://www.ibm.com/support/docview.wss?uid=ibm10882270)(<http://www.ibm.com/support/docview.wss?uid=ibm10882270>).

18. When the mailbox restore operation is finished, return to IBM Spectrum Protect Plus and click **Jobs and Operation > Active Resources > Databases**.
19. Select the resource, and then click **Cancel File Restore** to end the granular restore process.

Restoring Database Availability Group backups

With IBM Spectrum Protect Plus, you can restore a Microsoft Exchange Server Database Availability Group (DAG) backup to the original instance or to an alternative instance.

About this task

If you select the production restore method for the restore job, you must restore a replicated database copy to an active database copy. If you selected a passive database copy as the preferred target of backup operations, IBM Spectrum Protect Plus attempts to restore the database to this passive copy by default. The restore operation fails. In this situation, you can choose to restore the database to an alternative instance, and then select the active database copy.

If you select another restore method, the target location is a stand-alone recovery database that can be used on an active or passive DAG member.

Procedure

To define an Exchange restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. In the **Source select** page, complete the following steps:
 - a) Click the **View** menu and select **Database Availability Groups**.
 - b) In the **Availability Groups** list, click an Exchange instance to see the list of restore points for that instance and select the backup versions that you want to restore. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - c) Click the add to restore list icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the  icon next to the item.
 - d) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.

Option	Description
	<p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. In the **Restore method** page, choose from the following options:

- **Test.** Choose this option to restore the data from the vSnap repository directly. This restore type might be used for testing purposes.
- **Production.** Choose this option to restore the full database with a full-copy data restore operation. This restore operation is for permanent use of the restored database.

Click **Next** to continue.

6. In the **Set destination** page, specify where you want to restore the database and click **Next**.

Restore to original instance

Select this option to restore the database to the original server.

Restore to alternate instance

Select this option to restore the database to a local destination that is different from the original server, then select the alternative location from the list of available servers.



Attention: When you choose the destination, you must select an active node as the destination; otherwise, the restore operation fails.

- Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options

Choose from the following recovery options:

No Recovery

This option skips any rollforward recovery after the restore operation. The database remains in a Rollforward pending state until you decide whether you want to run the rollforward recovery manually.

Recover until end of backup

Restore the selected database to the state at the time the backup was created.

Recover until end of available logs

This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time

When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options

Set the application options:

Maximum Parallel Streams per Database

Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

- Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
- Take one of the following actions on the **Schedule** page:
 - If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
- On the **Review** page, review your restore job settings and click **Submit** to create the job.
The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.

Accessing Exchange database files with instant access mode

You can access the Exchange database files by using the instant access restore type and mount the database files from the vSnap volume to an application server.

About this task

In instant access mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery of data from the files in the vSnap volume.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > Exchange > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:

Option	Description
	<ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane. Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.

Option	Description
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Set destination** page, specify where you want to mount the database files and click **Next**.

Option	Description
Restore to original location	Select this option to mount the database files to the original server.
Restore to alternate location	Select this option to mount the database files to a local destination that is different from the original server, and then select the alternative location from the list of available servers.

6. On the **Restore Method** page, choose **Instant Access**, and then click **Next**.
7. Optional: On the **Job options** page, configure additional options if necessary and click **Next** to continue.
8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see [Configuring scripts](#). Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
- If you are running an on-demand job, click **Next**.
 - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.
The restore job is created, and you can check on its status in **Jobs and Operations > Running Jobs**.
11. You can now access the Exchange database files on the application server mount point, and carry out any Exchange related or custom actions you want to do.
- Note:** The Exchange database files on the mount point are read/write. However, updating them does not modify the original backup.
12. When you are finished with the instant access restore operation, click **Jobs and Operations > Active Resources > Databases**.

13. Select the resource, and then click **End Instant Disk Restore** to remove the mounted database and end the restore process.

MongoDB

After you successfully add MongoDB instances to IBM Spectrum Protect Plus, you can start to protect the data in your MongoDB databases. Create service level agreement (SLA) policies to back up and maintain MongoDB data.

Prerequisites for MongoDB

All system requirements and prerequisites for the IBM Spectrum Protect Plus MongoDB application server must be met before you start protecting MongoDB data with IBM Spectrum Protect Plus.

For MongoDB system requirements, see [MongoDB system requirements](#).

To meet the prerequisites for MongoDB, complete the following checks and actions.

1. Ensure you have met the space prerequisites, as described in [Space requirements for MongoDB protection](#).
2. Set the file size limit for the MongoDB instance user with the command **ulimit -f** to unlimited. Alternatively, set the value to sufficiently high to allow the copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the MongoDB instance to finalize the configuration.
3. If you are running MongoDB in a Linux environment, ensure that the installed sudo version is at a supported level.

For more information about the version level, see [“MongoDB requirements” on page 22](#). For information about setting sudo privileges, see [“Setting sudo privileges” on page 354](#).
4. If your MongoDB databases are protected by authentication, you must set up role-based access control. For more information, see [“Roles for MongoDB” on page 353](#).
5. Each MongoDB instance to be protected must be registered on IBM Spectrum Protect Plus. After the instances are registered, IBM Spectrum Protect Plus runs an inventory to detect MongoDB resources. Ensure that all instances that you want to protect are detected and listed correctly.
6. Ensure that the SSH service is running on port 22 on the server, and that firewalls are configured to allow IBM Spectrum Protect Plus to connect to the server with SSH. The SFTP subsystem for SSH must be enabled.
7. Ensure that you do not configure nested mount points.

Restrictions

The following restrictions apply to the MongoDB application server:

- MongoDB sharded cluster configurations are detected when you run an inventory, but these resources are not eligible for backup or restore operations.
- Unicode characters in MongoDB file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.

Virtualization

Protect your MongoDB environment with IBM Spectrum Protect Plus. The MongoDB application server that is installed on a VMware or Kernel-based Virtual Machine (KVM) virtual machine is protected when MongoDB is running on a supported operating system.

Roles for MongoDB

You must define role-based access control (RBAC) roles for the MongoDB agent users if authentication is enabled on the MongoDB database. When the roles are set up, users can protect and monitor MongoDB resources with IBM Spectrum Protect Plus in accordance with the users' defined roles.

Role-based access control for MongoDB

For each MongoDB user, specify access roles by using a command similar to the following example:

```
use admin
db.grantRolesToUser("<username>",
[ { role: "hostManager", db: "admin" },
  { role: "clusterManager", db: "admin" } ] )
```

The following roles are available:

hostManager

This role provides access to the **fsyncLock** command. This access is required for application-consistent backups of MongoDB databases where journaling is not enabled. This role also provides access to the shutdown command, which is used during a restore operation to shut down the MongoDB server instance that the restore is directed to.

clusterMonitor

This role provides access to commands for monitoring and reading the state of the MongoDB database. The following commands are available to users with this role:

- **getCmdLineOpts**
- **serverVersion**
- **replSetGetConfig**
- **replSetGetStatus**
- **isMaster**
- **listShards**

clusterManager

This role is only required only for running test restore operations of replica sets. Users who run the **replSetReconfig** command can create the restored instance of a single node replica set. This role enables read and write access during test restore operations of replica sets. Without this access, the node in the replica set would remain in the REMOVED state without read and write access. In addition, this role provides access to commands for reading the state of the MongoDB database. The following commands are available for this role:

- **replSetReconfig**
- **getCmdLineOpts**
- **serverVersion**
- **replSetGetConfig**
- **replSetGetStatus**
- **isMaster**
- **listShards**

Space prerequisites for MongoDB protection

Before you start backing up MongoDB data, ensure that you have enough free space on the target and source hosts, and in the vSnap repository. Extra space is required to store temporary Logical Volume Manager (LVM) backups of logical volumes where the MongoDB data is located. These temporary backups, that are known as LVM snapshots, are created automatically by the MongoDB agent.

LVM snapshots

LVM snapshots are point-in-time copies of LVM logical volumes. After the file copy operation finishes, earlier LVM snapshots are removed by the IBM Spectrum Protect Plus MongoDB agent in a cleanup operation.

For each LVM snapshot logical volume, you must allocate at least 10 percent free space in the volume group. If there is enough free space in the volume group, the IBM Spectrum Protect Plus MongoDB agent reserves up to 25 percent of the source logical volume size for the snapshot logical volume.

Linux LVM2

When you run a MongoDB backup operation, MongoDB requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system for each logical volume with data or logs for the selected database. On Linux systems, logical volumes are managed by LVM2.

A software-based LVM2 snapshot is taken as a new logical volume on the same volume group. The snapshot volumes are temporarily mounted on the same machine that runs the MongoDB instance so that they can be transferred to the vSnap repository.

On Linux, the LVM2 volume manager stores the snapshot of a logical volume within the same volume group. There must be enough space available to store the logical volume. The logical volume grows in size as the data changes on the source volume for the lifetime of the snapshot.

Setting sudo privileges

To use IBM Spectrum Protect Plus to protect your data, you must install the required version of the sudo program.

About this task

Set up a dedicated IBM Spectrum Protect Plus agent user with the required superuser privileges for sudo. This configuration enables agent users to run commands without a password.

Procedure

1. Create an agent user by issuing the following command:

```
useradd -m agent
```

where *agent* specifies the name of the IBM Spectrum Protect Plus agent user.

2. Set a password for the new user by issuing the following command:

```
passwd mongodb_agent
```

3. To enable superuser privileges for the agent user, set the `!requiretty` setting. At the end of the sudo configuration file, add the following lines:

```
Defaults:agent !requiretty
agent ALL=(ALL) NOPASSWD:ALL
```

Alternatively, if your sudoers file is configured to import configurations from another directory, for example `/etc/sudoers.d`, you can add the lines in the appropriate file in that directory.

Adding a MongoDB application server

To start protecting MongoDB resources, you must add the server that hosts your MongoDB instances, and set credentials for the instances. Repeat the procedure to add all the servers that host MongoDB resources.

About this task

To add a MongoDB application server to IBM Spectrum Protect Plus, you must have the host address of the machine.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > MongoDB**.
2. In the **MongoDB** window, click **Manage Application Servers**, and click **Add Application Server** to add the host machine.
3. In the **Application Properties** form, enter the **Host Address**.
4. Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

5. Choose to register the host with a user or an SSH key.

If you select **User**, you can choose to enter a new user and password, or an existing user. If you select **SSH Key**, select the SSH key from the menu.

Restriction: Any user that is specified must have sudo privileges set up.

The screenshot shows the 'Manage Application Servers' form in the MongoDB interface. The form is titled 'Application Properties' and contains the following fields and controls:

- Host Address:** A text input field containing 'metali.ca.ibm.com'.
- Get server key:** A blue button with a white border, highlighted by a blue rectangular selection box.
- Key type:** A text input field.
- Key fingerprint:** A text input field.
- Authentication:** Two radio buttons: 'User' (selected) and 'SSH Key'.
- Use existing user:** A checkbox that is currently unchecked.
- User ID:** A text input field containing 'domain\user'.
- Password:** A text input field containing 'Password'.
- Buttons:** 'Cancel' and 'Save' buttons at the bottom left.

Figure 24. Adding a MongoDB agent

6. Click **Get Instances** to detect and list the MongoDB instances that are available on the host server that you are adding.

Each MongoDB instance is listed with its connection host address, status, and an indication of whether it is configured.



Attention: If you register more than one application server for one replica set, the instance name that is displayed might change after each inventory, backup, or restore operation. The host name of the most recently added application server that belongs to the replica set is used as part of the instance name. An inventory operation is run as part of backup and restore operations.

7. If you are using access control, configure an instance by setting credentials. Click **Set Credential**, and set the user ID, and password. Alternatively, you can select to use an existing user profile.

For more information about access control, see [Chapter 16, “Managing user access,”](#) on page 455.

When you set credentials, you assign MongoDB user roles for the backup and restore operations with access to role-protected MongoDB servers by using Salted Challenge Response Authentication Mechanism (SCRAM), or Challenge and response authentication. The MongoDB user that is assigned for the role-protected MongoDB server requires one of the following access levels to protect resources:

- *Host Manager:* manages the database as the administrator. This role is required for taking and managing snapshots.
- *Cluster Administrator:* retrieves configuration information and runs test mode restore operations of MongoDB replica sets. This role is required to reconfigure test mode restore operations of MongoDB replica sets for data queries.
- *Cluster Monitor:* monitors the protection of MongoDB resources, and retrieves configuration information.

- Optional: Set the option **Maximum concurrent databases** by entering a number in the field.
- Save the form, and repeat the steps to add other MongoDB application servers to IBM Spectrum Protect Plus.

What to do next

After you add MongoDB application servers to IBM Spectrum Protect Plus, an inventory is automatically run on each application server to detect the relevant databases in those instances.

To verify that the databases are added, review the job log. Go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as a wildcard in the name.

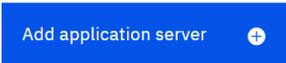
Databases must be detected to ensure that they can be protected. For instructions about running a manual inventory, see [Detecting MongoDB resources](#).

Registering a MongoDB Ops Manager Application Database for protection

To protect your MongoDB Ops Manager Application Database, you must first register the Ops Manager host address with IBM Spectrum Protect Plus.

Procedure

- In the navigation panel, expand **Manage Protection > Databases > MongoDB**.
- In the **MongoDB** window, click **Manage Application Servers**, and click **Add Application Server**.

A blue rectangular button with the text "Add application server" and a white plus sign icon to its right.

- In the Application Properties form, enter the host address for the Ops Manager Application Database. Get instances and set credentials by following the steps outlined in [“Adding a MongoDB application server”](#) on page 355.

The Ops Manager Application Database is listed in the Instances table as shown in the following example:

```
metali8.limerick.ie.ibm.com Connection: '333.0.5.1:88888' Ops Manager Application Database
```

What to do next

The MongoDB Ops Manager Application Database is available for backing up. You can define backup and restore jobs to protect your data. To regularly back up your data, define a backup job that includes a service level agreement (SLA) policy. For more information, see [“Backing up MongoDB data”](#) on page 359 and [“Defining a regular service level agreement job”](#) on page 360.

Detecting MongoDB resources

After you add your MongoDB application servers to IBM Spectrum Protect Plus, an inventory is run automatically to detect all MongoDB instances and databases. You can run a manual inventory on any application server to detect, list, and store all MongoDB databases for the selected host.

Before you begin

Ensure that you added your MongoDB application servers to IBM Spectrum Protect Plus. For instructions, see [Adding a MongoDB application server](#).

Important:

Enable IP hostname resolution for all replica set member hosts on any MongoDB server that is registered in IBM Spectrum Protect Plus. When a MongoDB replica set member IP address is not able to be resolved to a hostname, the inventory will then use the IP address as the key property for the *clusterHosts* property which includes dots. Cataloging of the database instance fails on inventory with the following error message:

```
Cataloging failed for server <servername>:  
org.springframework.data.mapping.model.MappingException:  
Map key <hostname> contains dots but no replacement was configured!  
Make sure map keys don't contain dots in the first place or configure an appropriate  
replacement!
```

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > MongoDB**.

Tip: To add more MongoDB instances to the **Instances** pane, follow the instructions in [Adding a MongoDB application server](#).

2. Click **Run Inventory**.

When the inventory is running, the button changes to **Inventory In Progress**. You can run an inventory on any available application servers, but you can run only one inventory process at a time.

To monitor the inventory job, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

3. Click an instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your MongoDB application server and rerun the inventory. In some cases, certain databases are marked as ineligible for backup; hover over the database to reveal the reason why.

Tip: To return to the list of instances, click the **Instances** link in the **Backup MongoDB** pane.



Attention: If you register more than one application server for one replica set, the instance name that is displayed might change after each inventory, backup, or restore operation. The host name of the most recently inventoried application server that belongs to the replica set is used as part of the instance name. An inventory operation is run as part of backup and restore operations.

What to do next

To start protecting MongoDB databases that are cataloged in the selected instance, apply a service level agreement (SLA) policy to the instance. For instructions about setting an SLA policy, see [Defining an SLA policy](#).

Testing the MongoDB connection

After you add a MongoDB application server, you can test the connection. The test verifies communication between IBM Spectrum Protect Plus and the MongoDB server. It also checks that the correct sudo permissions area available for the user who is running the test.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > MongoDB**.
2. In the **MongoDB** window, click **Manage Application Servers**, and select the host address that you want to test.

A list of the MongoDB application servers that are available is shown.

3. Click **Actions** and choose **Test** to start the verification tests for physical and remote system connections and settings.

The test report displays a list that includes tests for the physical host network configuration, and tests for the remote server installation on the host.

4. Click **OK** to close the test report. If issues are reported, fix the issues and rerun the test to verify the fixes.

Backing up MongoDB data

You can define backup jobs to protect your MongoDB data. To regularly back up your data, define a backup job that includes a service level agreement (SLA) policy.

Before you begin

During the initial backup operation, IBM Spectrum Protect Plus creates a vSnap volume and NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus MongoDB agent mounts the share on the MongoDB server where the backup is completed.

Review the following prerequisites before you create a backup job definition:

- Add the application servers that you want to back up. For the procedure, see [Adding a MongoDB application server](#).
- Configure an SLA Policy. For the procedure, see [Defining a Service Level Agreement backup job](#).
- Before an IBM Spectrum Protect Plus user can set up backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources, and backup and restore operations, in the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,” on page 455](#) and [“Roles for MongoDB” on page 353](#).

Restriction: Do not run inventory jobs at the same time that backup jobs are scheduled.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > MongoDB**.
2. Select the check box for the instance that you want to back up.

Under each MongoDB instance, data to be backed up is listed as **ALL**. Each instance in the Instances pane is listed by instance name, version, and the applied SLA policy.

3. Click **Select Options** to specify the number of parallel streams for the backup operation, and then click **Save**. By selecting an appropriate number of parallel streams, you can minimize the time that is required for the backup job.

The saved options are used for all backup jobs for this instance as selected.

4. To run the backup job with these options, click the instance name, select the **ALL** database representation, and click **Run**.

The backup job begins, and you can view the details in **Jobs and Operations > Running Jobs**.

Tip: The **Run** button is only enabled if an SLA policy is applied to the **ALL** representation of the databases.

To run an on-demand backup job for multiple databases that are associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job” on page 439](#).

5. Select the instance again, and click **Select an SLA Policy** to choose an SLA policy.
6. Save the SLA selection.

To define a new SLA or to edit an existing policy with custom retention and frequency rates, select **Manage Protection > Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define policy preferences.

What to do next

After the SLA policy is saved, you can run the policy at any time by clicking **Actions** for that policy name, and selecting **Start**. The status in the log changes to show that the backup job is in the Running state.

To cancel a job that is running, click **Actions** for that policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed up data, or **No** to discard the backup.

Important: If the backup operation fails with an error, follow the procedure described in [“Troubleshooting failed backup operations for large Db2, MongoDB, and SAP HANA databases”](#) on page 473.

Defining a regular service level agreement job

After your MongoDB instances are listed, select and apply an SLA policy to start protecting your data.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > MongoDB**.
2. Select the MongoDB instance to back up all the data in that instance.
3. Click **Select an SLA Policy** to select an SLA policy, and then click **Save**.
Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy or edit an existing policy. For more information see [“Creating an SLA policy for databases and file systems”](#) on page 206.
4. Optional: To enable multiple backup streams to reduce the time that is taken to back up large databases, click **Select Options** and enter a number of parallel streams. Save your changes.

Tips for configuring the option:

Review the following tips to help you configure the option for the backup job:

- To set the option for child resources to the same value as the parent, click **Set all options to inherit**.
 - If multiple resources were selected for the backup job, the option value is indeterminate. If you change a value for the option, that value is used for all selected resources after you click **Save**.
 - If an option value is changed from the previously saved value, the option is shown in yellow.
 - To close the **Options** pane without saving changes, click **Select Options**.
5. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.
For more information about SLA configuration options, see [“Setting SLA configuration options for your backup”](#) on page 361.
 6. To run the policy outside of the scheduled job, select the instance. Click the **Actions** button and select **Start**. The status changes to **Running** for your chosen SLA and you can follow the progress of the job in the log shown.

What to do next

After the SLA policy is saved, you can run the policy at any time by clicking **Actions** for that policy name, and selecting **Start**. The status in the log changes to show that the backup job is in the Running state.

To cancel a job that is running, click **Actions** for that policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed up data, or **No** to discard the backup.

Important: If the backup operation fails with an error, follow the procedure described in [“Troubleshooting failed backup operations for large Db2, MongoDB, and SAP HANA databases”](#) on page 473.

Setting SLA configuration options for your backup

After you set up a service level agreement (SLA) policy for your backup job, you can choose to configure extra options for that job. Additional SLA options include running scripts, and forcing a full base backup.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job that you are configuring, click the clipboard icon  to specify additional configuration options.
If the job is already configured, click on the icon to edit the configuration.
2. Click **Pre-Script** and define the prescript configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.
3. Click **Post-Script** and define the PostScript configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.

Scripts and script servers are configured on the **System Configuration > Script** page. For more information about working with scripts, see **Configuring scripts**.

4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.
If this option is selected, the backup or restore operation is reattempted after an initial fail, and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.
5. Skip **Exclude Resources** for MongoDB SLA options, as you cannot specify resources to exclude. Instances are backed up rather than individual databases.
6. To create a full, new backup of a MongoDB instance, select **Force full backup of resources**.
A full new backup of that resource is created to replace the existing backup of that resource for one occurrence only. After that the resource is backed up incrementally as before.

Restoring MongoDB data

To restore data, define a job that restores data to the latest backup or select an earlier backup copy. Choose to restore data to the original instance or to an alternative instance on a different machine, creating a cloned copy. Define and save the restore job to run as an ad hoc operation, or to run regularly as a scheduled job.

Before you begin

Before you create a restore job for MongoDB, ensure that the following requirements are met:

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up MongoDB data” on page 359](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#), and [“Roles for MongoDB” on page 353](#).
- Enough disk space is allocated at the target server for the restore operation.
- Dedicated volumes are allocated for file copying.
- The same directory structure and layout are available on both the target and source servers.

- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

For restore operations to alternative instances, MongoDB must be at the same version level on the target and host machines.

For more information about space requirements, see [Space prerequisites for MongoDB protection](#). For more information about prerequisites and setup, see [Prerequisites for MongoDB](#).

Procedure

To define a MongoDB restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > MongoDB > Create job**, and then select **Restore** to open the Restore wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > MongoDB**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the add to restore list icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:

Option	Description
	<ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane. Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.

Option	Description
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, choose the type of restore operation, and click **Next** to continue.

- **Test:** In this mode, the agent creates a database by using the data files directly from the vSnap repository. This option is available only when you are restoring data to an alternative instance. Members of replica sets will not be reconfigured after the MongoDB server is started. The server is started as a single-node replica set.
- **Production:** In this mode, the MongoDB application server first copies the files from the vSnap repository to the target host. The copied data is then used to start the database. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.
- **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery from the files in the vSnap repository.

For test mode or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

6. On the **Set destination** page, select **Restore to original instance** to restore to the original server, or **Restore to alternate instance** to restore to a different location that you can select from the locations listed.

For more information about restoring data to the original instance, see [Restoring to the original instance](#). For more information about restoring your data to an alternative instance, see [Restoring to an alternate instance](#).

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

In the **Recovery Options** section, the **Recover until end of backup** for MongoDB is selected by default. This option recovers the selected data to the state it was in at the time the backup was created. The recovery operation makes use of the log files that are included in the MongoDB backup.

Application Options

Set the application options:

Overwrite existing database

Enable this option to allow the restore job to overwrite the selected database. If this option is not selected, the restore job fails when data with the same name is found during the restore process.



Attention: Ensure that no other data shares the same local database directory as the original data or the data will be overwritten.

Maximum Parallel Streams per Database

Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might speed up restore operations, but high bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring a MongoDB database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace existing databases with the same name during a restore operation. During an instant disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

Continue with restores of other selected databases even if one fails

If one database in the instance is not successfully restored, the restore operation continues for all other data that is being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

Mount Point Prefix

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

- Optional: On the **Apply scripts** page, specify scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported on Windows operating systems while shell scripts are supported on Linux operating systems.

Pre-Script

Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration > Script**.

Post-Script

Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration > Script** page.

Continue job/task on script error

Select this option to continue running the job when the script that is associated with the job fails. When this option is enabled, in the event that a script completes processing with a nonzero return code, the backup or restore job continues to run and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task is reported as FAILED.

Click **Next** to continue.

9. On the **Schedule** page, click **Next** to start on-demand jobs after you complete the Restore wizard. For recurring jobs, enter a name for the job schedule, and specify how often and when to start the restore job.
10. On the **Review** page, review your restore job settings.



Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

11. To proceed with the job, click **Submit**. To cancel the job, navigate to **Jobs and Operations** and click the **Schedule** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Results

A few moments after you select **Restore**, the **onDemandRestore** job is added to the **Jobs and Operations > Running Jobs** pane. Click the record to show the step-by-step details of the operation. You can also download the zipped log file by clicking **Download.zip**. For any other jobs, click the **Running Jobs** or **Job History** tabs and click the job to display its details.

The IP address and port for the restored server can be found in the log file for the restore operation. Navigate to **Jobs and Operations > Running Jobs** to find the logs for your restore operation.

For information about restoring data to the original instance, see [Restoring to the original instance](#). For information about restoring your data to an alternative instance, see [Restoring to an alternate instance](#).

Restoring MongoDB data to the original instance

You can restore a MongoDB instance to the original host and choose between restoring to the latest backup or an earlier MongoDB database backup version. When you restore data to its original instance, you cannot rename it. This restore option runs a full production restoration of data, and existing data is overwritten at the target site if the **Overwrite existing databases** application option is selected.

Before you begin

Before you create a restore job for MongoDB, ensure that the following requirements are met:

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up MongoDB data”](#) on page 359.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see [Chapter 16, “Managing user access,”](#) on page 455, and [“Roles for MongoDB”](#) on page 353.
- Enough disk space is allocated at the target server for the restore operation.
- Dedicated volumes are allocated for file copying.
- The same directory structure and layout are available on both the target and source servers.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

For more information about space requirements, see [Space prerequisites for MongoDB protection](#). For more information about prerequisites and setup, see [Prerequisites for MongoDB](#).

Procedure

1. In the navigation panel, click **Manage Protection > Databases > MongoDB > Create job**, and then select **Restore** to open the Restore wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > MongoDB**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the add to restore list icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage

Option	Description
	types Backup , Replication , and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.

Option	Description
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, choose the type of restore operation, and click **Next** to continue.

- **Production**

To recover an entire instance to the original instance, the preferred method is to choose this option with the overwrite application option. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.

- **Test**

Choose this option to restore data to the same server but using a different port.

- **Instant Access**

Choose this option to mount the backup to the application server without restoring the data or overwriting the data.

Click **Next** to continue.

For test mode or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

6. On the **Set destination** page, choose **Restore to original instance** and click **Next**.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

In the **Recovery Options** section, the **Recover until end of backup** for MongoDB is selected by default. This option recovers the selected data to the state it was in at the time the backup was created. The recovery operation makes use of the log files that are included in the MongoDB backup.

Application Options

Set the application options:

Overwrite existing database

Enable this option to allow the restore job to overwrite the selected database. If this option is not selected, the restore job fails when data with the same name is found during the restore process.



Attention: Ensure that no other data shares the same local database directory as the original data or the data will be overwritten.

Maximum Parallel Streams per Database

Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might speed up restore operations, but high bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring a MongoDB database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace existing databases with the same name during a restore operation. During an instant disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

Continue with restores of other selected databases even if one fails

If one database in the instance is not successfully restored, the restore operation continues for all other data that is being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

Mount Point Prefix

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

- Optional: On the **Apply scripts** page, specify scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported on Windows operating systems while shell scripts are supported on Linux operating systems.

Pre-Script

Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration > Script**.

Post-Script

Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration > Script** page.

Continue job/task on script error

Select this option to continue running the job when the script that is associated with the job fails. When this option is enabled, in the event that a script completes processing with a nonzero return code, the backup or restore job continues to run and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task is reported as FAILED.

Click **Next** to continue.

- On the **Schedule** page, click **Next** to start on-demand jobs after you complete the Restore wizard. For recurring jobs, enter a name for the job schedule, and specify how often and when to start the restore job.
- On the **Review** page, review your restore job settings.



Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

- To proceed with the job, click **Submit**. To cancel the job, navigate to **Jobs and Operations** and click the **Schedule** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Restoring MongoDB data to an alternative instance

You can select a MongoDB database backup and restore it to an alternative host. You can also choose to restore a database to a different vSnap repository, or you can rename the database. This process creates an exact copy of the instance on a different host.

Before you begin

Before you create a restore job for MongoDB, ensure that the following requirements are met:

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up MongoDB data” on page 359](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#), and [“Roles for MongoDB” on page 353](#).
- Enough disk space is allocated at the target server for the restore operation.
- Dedicated volumes are allocated for file copying.
- The same directory structure and layout are available on both the target and source servers.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

For restore operations to alternative instances, MongoDB must be at the same version level on the target and host machines.

For more information about space requirements, see [Space prerequisites for MongoDB protection](#). For more information about prerequisites and setup, see [Prerequisites for MongoDB](#).

Procedure

1. In the navigation panel, click **Manage Protection > Databases > MongoDB > Create job**, and then select **Restore** to open the Restore wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > MongoDB**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the add to restore list icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	Select a type of location from which to restore data:

Option	Description
	<p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	<p>For on-demand restore operations, specify a range of dates to show the available snapshots within that range.</p>
Restore Point	<p>For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.</p>
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

5. On the **Restore method** page, choose the type of restore operation, and click **Next** to continue.

- **Test:** In this mode, the agent creates a database by using the data files directly from the vSnap repository. This option is available only when you are restoring data to an alternative instance. Members of replica sets will not be reconfigured after the MongoDB server is started. The server is started as a single-node replica set.
- **Production:** In this mode, the MongoDB application server first copies the files from the vSnap repository to the target host. The copied data is then used to start the database. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.

- **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery from the files in the vSnap repository.

For test mode or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

6. In the **Set destination** page, choose **Restore to alternate instance** and select the target instance that you want to restore the data to.

The original instance is not selectable because you cannot overwrite the original data when you select **Restore to alternate instance**. You also cannot select instances on different versions levels or instances on the same host as the original instance.

Click **Next** to continue.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

In the **Recovery Options** section, the **Recover until end of backup** for MongoDB is selected by default. This option recovers the selected data to the state it was in at the time the backup was created. The recovery operation makes use of the log files that are included in the MongoDB backup.

Application Options

Set the application options:

Overwrite existing database

Enable this option to allow the restore job to overwrite the selected database. If this option is not selected, the restore job fails when data with the same name is found during the restore process.



Attention: Ensure that no other data shares the same local database directory as the original data or the data will be overwritten.

Maximum Parallel Streams per Database

Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might speed up restore operations, but high bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring a MongoDB database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace existing databases with the same name during a restore operation. During an instant disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

Continue with restores of other selected databases even if one fails

If one database in the instance is not successfully restored, the restore operation continues for all other data that is being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

Mount Point Prefix

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

- Optional: On the **Apply scripts** page, specify scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported on Windows operating systems while shell scripts are supported on Linux operating systems.

Pre-Script

Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration > Script**.

Post-Script

Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration > Script** page.

Continue job/task on script error

Select this option to continue running the job when the script that is associated with the job fails. When this option is enabled, in the event that a script completes processing with a nonzero return code, the backup or restore job continues to run and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task is reported as FAILED.

Click **Next** to continue.

- On the **Schedule** page, click **Next** to start on-demand jobs after you complete the Restore wizard. For recurring jobs, enter a name for the job schedule, and specify how often and when to start the restore job.
- On the **Review** page, review your restore job settings.



Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

- To proceed with the job, click **Submit**. To cancel the job, navigate to **Jobs and Operations** and click the **Schedule** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Using a granular restore operation for MongoDB

You can restore specific MongoDB databases or collections by using a granular restore operation. For a granular restore operation, first run a test restore job and then run the appropriate MongoDB commands.

Before you begin

If authentication is enabled, you must provide credentials for users so that they can correct permissions on the instance in the test restore operation.

About this task

The granular restore operation for MongoDB is based on a test mode restore job. When you run the test restore job on IBM Spectrum Protect Plus, and the **mongodump** and **mongorestore** commands on the MongoDB server, you can access individual databases or collections from the recovery source.

Use this procedure to complete either of the following tasks:

- Restore any number of databases by using the **mongodump** and **mongorestore** commands for the databases that you require.
- Restore any number of collections by using the **mongodump** and **mongorestore** commands for the collections that you require.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > MongoDB > Create job**, and then select **Restore** to open the **Restore** wizard.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the add to restore list icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon  next to the item.
 - c) Click **Next** to continue.
3. On the **Restore method** page, select **Test**, and click **Next** to continue with the test restore process.
4. On the **Set destination** page, choose **Restore to alternate instance**, and select the target instance that you want to restore the data to.

You cannot select the original instance is not selectable as you cannot overwrite the original data when you select **Restore to alternate instance**. Instances on different versions levels cannot be selected. Other instances on the same host as the original instance, cannot be selected either.

Click **Next** to continue.

5. Proceed through the restore wizard pages and select the required options.
6. On the **Review** page, review your restore job settings.



Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

7. Log on to the MongoDB server to which the test restore job is directed.
8. Run the MongoDB system command `ps -ef | grep mongod` to find the temporary recovery MongoDB instance location.
9. Run the MongoDB `mongodump` command to create a dump file of any specific database or collection.

Use the appropriate command. The first command is for a database and the second command is for a collection:

```
mongodump --host <hostname> --port <port> --db <dbname> <dumpfolder>
```

Or,

```
mongodump --host <hostname> --port <port> --collection <collectionname> <dumpfolder>
```

10. Run the **mongorestore** command to restore the dump file into any MongoDB instance. Choose either the original MongoDB instance that the backup was created for, or any alternative instance.

Use the appropriate command. The first command is for a database and the second command is for a collection:

```
mongorestore --host <hostname> --port <port> --db <dbname> <dumpfolder>\<dbname>
```

Or,

```
mongorestore --host <hostname> --port <port> --collection <collectionname>  
<dumpfolder>\<dbname>
```

11. When the database or collection restore operation finishes, go to **Jobs and Operations > Active Resources > Databases**.

12. Click **Cancel Restore** to end the granular restore procedure.

Backing up and restoring Oracle data

To protect Oracle content, first register the Oracle instance so that IBM Spectrum Protect Plus recognizes it. Then create jobs for backup and restore operations.

Ensure that your Oracle environment meets the system requirements in [“Oracle Server database backup and restore requirements”](#) on page 22.

Adding an Oracle application server

When an Oracle application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Procedure

To register an Oracle application server, complete the following steps.

1. In the navigation panel, click **Manage Protection > Databases > Oracle**.
2. Click **Manage Application Servers**.
3. Click **Add Application Server** to add the host machine.
4. In the **Application Properties** pane, enter the **Host Address**.
The host address is a resolvable IP address, or a resolvable path and machine name.
5. Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

When upgrading to the IBM Spectrum Protect Plus latest version, systems that are already registered in the previous version are set to trust on first use (TOFU) and the SSH key fingerprint will automatically be added to the registration information in the catalog.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

6. Select **User** or **SSH key**.

Option	Description
User	Click this option to specify an existing user or enter a user ID and password. The user must have sudo privileges set up. Populate the fields as follows: Use existing user Select this check box to use a previously entered user name and password for the application server. Select a user name from the Select user list. UserID Enter your user name for the application server. If the virtual machine is attached to a domain, the user identity follows the default <i>domain\name</i> format. If the user is a local administrator, use the <i>local_administrator</i> format.

Option	Description
	Password Enter your password for the application server.
SSH Key	Click this option to use an SSH key. Select a key from the Select a SSH key list.

7. To protect multithreaded databases in Oracle 12c and later versions, provide credentials for the databases:
 - a) Click **Get databases** to detect and list the Oracle databases on the host server that you are adding. Each Oracle database is listed with its name, status, and an indication of whether credentials were previously specified for the database.
 - b) For each multithreaded database that you want to protect, click **Set Credential** and specify the user ID and password. Alternatively, you can select an existing user from the **Select user** list. You must specify the credentials of an Oracle database user who has SYSDBA privileges.
8. In **Maximum concurrent databases**, set the maximum number of databases to back up concurrently on the server.
Server performance is impacted when many databases are backed up concurrently, as each database utilizes multiple threads and consumes bandwidth when copying data. Use this option to control the impact on server resources and minimize the impact on production operations.
9. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the application server to the IBM Spectrum Protect Plus database, and then catalogs the instance.
If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

What to do next

After you add the Oracle application server, complete the following action:

Action	How to
Assign user permissions to the application server.	See “Creating a role” on page 463 .

Related concepts

[“Managing user access” on page 455](#)

By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Related tasks

[“Backing up Oracle data” on page 379](#)

Use a backup job to back up Oracle environments with snapshots.

[“Restoring Oracle data” on page 382](#)

Use a restore job to restore an Oracle environment from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version that is selected during the job definition creation and creates a Network Files System (NFS) share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle server where the restore job is to be run. For Oracle Real Application Clusters (RAC), the restore job is run on all nodes in the cluster.

Detecting Oracle resources

Oracle resources are automatically detected after the application server is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the application server was added.

Before you begin

IBM Spectrum Protect Plus uses the Oracle instance inventory files that are named `inventory.xml` to locate databases. These files are contained in a `ContentsXML` subdirectory for each instance that is pointed to by the `/etc/oraInst.loc` file. Ensure that the inventory files and the `oraInst.loc` file are updated with the latest changes to your environment. If the inventory files are outdated or corrupted, or if the `oraInst.loc` file does not point to the inventory locations, the associated databases are not be recognized in IBM Spectrum Protect Plus. For more information about managing the Oracle instance inventory, see https://docs.oracle.com/cd/E11882_01/em.112/e12255/oui2_manage_oracle_homes.htm#OUICG143.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Oracle**.
2. In the list of Oracle instances, select an instance or click the link for the instance to navigate to the resource that you want. For example, if you want to run an inventory job for an individual database in the instance, click the instance link and then select a virtual machine.
3. Click **Run Inventory**.

Testing connection to an Oracle application server

You can test the connection to an Oracle host. The test function verifies communication with the host and tests DNS settings between the IBM Spectrum Protect Plus virtual appliance and the host.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Oracle**.
2. Click **Manage Application Servers**.
3. In the list of hosts, click **Test** in the **Actions** menu for the host.

Backing up Oracle data

Use a backup job to back up Oracle environments with snapshots.

Before you begin

Review the following information:

- To ensure that file system permissions are retained correctly when IBM Spectrum Protect Plus moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (for example, `oracle`, `oinstall`, `dba`) are consistent across all the servers. Refer to Oracle documentation for recommended `uid` and `gid` values.
- If an Oracle Inventory job runs at the same time or short period after an Oracle backup job, copy errors might occur because of temporary mounts that are created during the backup job. As a best practice, schedule Oracle Inventory jobs so that they do not overlap with Oracle backup jobs.
- Avoid configuring log backup for a single Oracle database by using multiple backup jobs. If a single Oracle database is added to multiple job definitions with log backup enabled, a log backup from one job could truncate a log before it is backed up by the next job. This might cause point-in-time restore jobs to fail.

- Avoid scheduling log backups at the same time as an SLA backup job for the same Oracle database. If a log backup occurs at the same time as the backup task of an SLA backup, the SLA backup job may fail. Additionally, ad-hoc backups should not be started if they will run at the same time as scheduled log backups.
- Point-in-time recovery is not supported when one or more data files are added to the database in the period between the chosen point-in-time and the time that the preceding backup job ran.

Take the following actions:

- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,” on page 455](#).
- Register the providers that you want to back up. For more information, see [“Adding an Oracle application server” on page 377](#).
- Configure SLA policies. For more information, see [“Create backup policies” on page 98](#).

About this task

During the initial base backup, IBM Spectrum Protect Plus creates a vSnap volume and an NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus agent mounts the share on the Oracle server where the backup is to be completed.

In the case of Oracle Real Application Clusters (RAC), the backup is completed from any one node in the cluster. When the backup job is completed, the IBM Spectrum Protect Plus agent unmounts the share from the Oracle server and creates a vSnap snapshot of the backup volume.

IBM Spectrum Protect Plus can protect multithreaded databases in Oracle 12c and later versions. For instructions about enabling IBM Spectrum Protect Plus to protect multithreaded databases, see [“Adding an Oracle application server” on page 377](#).

Procedure

To define an Oracle backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Oracle**.
2. Select Oracle homes, databases, and ASM diskgroups to back up. Use the search function to search for available instances.
3. Click **Select an SLA Policy** to add one or more SLA policies that meet your backup data criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

The job runs as defined by the SLA policies that you selected. To run the job manually, click **Jobs and Operations > Schedule**. Select the job and click **Actions > Start**.

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
 - To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job” on page 439](#).
5. To edit options before you create the job definition, click **Select Options**. Set the job definition options.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.

- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

Enable Log Backup

Enable Log Backup must be selected to allow for Oracle point-in-time restore.

Select **Enable Log Backup** to permit IBM Spectrum Protect Plus to automatically create a log backup volume and mount it to the application server. IBM Spectrum Protect Plus then automatically discovers the location of the existing primary archived log and uses cron to configure a scheduled job. The scheduled job completes a transaction log backup from the primary location to that log backup volume at the frequency specified through the **Frequency** setting.

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

The **Frequency** can be set to a value independent of the database backup frequency specified in the SLA Policy settings. For example, the SLA Policy may be configured to back up the database once per day while the log backup frequency could be set to once per 30 minutes.

For Oracle RAC, IBM Spectrum Protect Plus mounts the volume and configures the cron job on each of the cluster nodes. When the schedule is triggered, the jobs internally coordinate to ensure that any one active node completes the log backup and the other nodes take no action.

IBM Spectrum Protect Plus automatically manages the retention of logs in its own log backup volume based on the retention settings in the SLA policy.

Choose an option for **Truncate source logs after successful backup**. Select **Never**, the default, if you do not want to truncate source logs after a successful backup job. If this option is selected, archived logs on the primary log destination are not deleted, and Database Administrators must continue to manage those logs using their existing log retention policies. Select **Older than a specified number of days**, **Older than a specified number of hours**, or **Immediately after log backup** if you do want source logs to be truncated after successful backup. Selecting one of these options will automatically delete older archived logs from the database's primary archived log location after a number of days, hours or at the end of every successful database backup.

When the option **Truncate source logs after successful backup** is set to **Older than a specified number of days** or **Older than a specified number of hours**, set the retention of primary logs by entering a day or hour value in the **Primary log retention in days** or **Primary log retention in hours** field. This setting controls the quantity of archived logs that are retained in the primary archived log locations. For example, if **Older than a specified number of days** is selected and **Primary log retention in days** is set to **3**, IBM Spectrum Protect Plus deletes all archived logs older than three days from the primary archived log location at the end of every successful database backup.

Note: Users that upgrade from a previous version of IBM Spectrum Protect Plus that did not select the **Truncate source logs after successful backup** will have that option set to **Never** after upgrading. Similarly, users that opted to select the **Truncate source logs after successful backup** option and entered a number of days for **Primary log retention in days** will have the option set for **Older than a specified number of days** and the associated number of days set in **Primary log retention in days** field after upgrading.

Maximum Parallel Streams per Database

Set the maximum data stream per database to the backup storage. This setting applies to each database in the job definition. Multiple databases can be backed up in parallel if the value of the option is set to **1**. Multiple parallel streams might improve backup speed, but high bandwidth consumption might affect overall system performance.

6. When you are satisfied that the job-specific information is correct, click **Save**.

- To configure additional options, click the **Policy Options** clipboard icon  that is associated with the job in the **SLA Policy Status** section. Set the following additional policy options:

Pre-scripts and Post-scripts

Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backup operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

What to do next

After you create the backup job definition, complete the following action:

Action	How to
Create an Oracle Restore job definition.	See “Restoring Oracle data” on page 382 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Restoring Oracle data

Use a restore job to restore an Oracle environment from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version that is selected during the job definition creation and creates a Network Files System (NFS) share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle server where the restore job is to be run. For Oracle Real Application Clusters (RAC), the restore job is run on all nodes in the cluster.

Before you begin

Complete the following prerequisites:

- Create and run an Oracle backup job. For instructions, see [“Backing up Oracle data” on page 379](#).
- Before an IBM Spectrum Protect Plus user can restore data, the appropriate roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations by using the **Accounts** pane. For instructions, see [Chapter 16, “Managing user access,” on page 455](#).

Review the following restrictions:

- Point-in-time recovery is not supported if one or more data files were added to the database in the period between the chosen point in time and the time that the preceding backup job ran.
- If an Oracle database is mounted but not opened during a backup job, IBM Spectrum Protect Plus cannot determine the database **tempfile** settings that are related to **autoextensibility** and maximum size. When a database is restored from this restore point, IBM Spectrum Protect Plus cannot re-create the **tempfiles** with the original settings because they are unknown. Instead, **tempfiles** are created with default settings, `AUTOEXTEND ON` and `MAXSIZE 32767M`. After the restore job is completed, you can manually update the settings.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

About this task

The following restore modes are supported:

Instant access mode

In instant access mode, no further action is taken after mounting the share. Users can complete any custom recovery by using the files in the vSnap volume.

Test mode

In test mode, the agent creates a new database by using the data files directly from the vSnap volume.

Production mode

In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

Procedure

To define an Oracle restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > Oracle > Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > Oracle**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

- c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	<p>For on-demand restore operations, specify a range of dates to show the available snapshots within that range.</p>
Restore Point	<p>For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.</p>
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

- On the **Restore method** page, set the restore job to run in test, production, or instant access mode by default.

For test or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

Click **Next** to continue.

After the job is created, it can be run in test, production, or instant access mode in the **Job Sessions** pane.

6. On the **Set destination** page, specify where you want to restore the database and click **Next**.

Restore to original location

Select this option to restore the database to the original server.

Restore to alternate location

Select this option to restore the database to a local destination that is different from the original server, and then select the alternative location from the list of available servers.

7. On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options

Set the following point-in-time recovery options:

Recover until end of backup

Restore the selected database to the state at the time that the backup was created.

Recover until specific point in time

When log backup is enabled by using an Oracle Backup job definition, point-in-time restore options will be available when you create an Oracle Restore job definition. Select one of the following options, and then click **Save**:

- **By Time.** Select this option to configure a point-in-time recovery from a specific date and time.
- **By SCN.** Select this option to configure a point-in-time recovery by System Change Number (SCN).

IBM Spectrum Protect Plus finds the restore points that directly proceed and follow the selected point in time. During the recovery, the older data backup volume and the newer log backup volume are mounted. If the point in time occurred after the last backup, a temporary restore point is created.

Application Options

Set the application options:

Overwrite existing database

Enable this option to allow the restore job to overwrite the selected database. By default, this option is not selected.

Maximum Parallel Streams per Database

Set the maximum number of parallel data stream from the backup storage per database. This setting applies to each database in the job definition. If the value of the option is set to 1, multiple databases can still be restored in parallel. Multiple parallel streams might improve restore speed, but high bandwidth consumption might affect overall system performance.

Init Params

This option controls the initialization parameters that are used to start the recovered database in Oracle test and production workflows.

Source. This option is the default. IBM Spectrum Protect Plus uses the same initialization parameters as the source database, but with the following changes:

- Parameters that contain paths such as **control_files**, **db_recovery_file_dest**, or **log_archive_dest_*** are updated to reflect the new paths based on the renamed mount points of the recovered volumes.
- Parameters such as **audit_file_dest** and **diagnostic_dest** are updated to point to the appropriate location under the Oracle base directory on the destination server if the path differs from the source server.
- If a new name is specified for the database, the **db_name** and **db_unique_name** parameters are updated to reflect the new name.

- Cluster-related parameters such as **instance_number**, **thread**, and **cluster_database** are set automatically by IBM Spectrum Protect Plus, depending on the appropriate values for the destination.

Target. Customize the initialization parameters by specifying a template file that contains the initialization parameters that are used by IBM Spectrum Protect Plus.

The specified path must point to a plain text file that exists on the destination server and is readable by the IBM Spectrum Protect Plus user. The file must be in Oracle pfile format, consisting of lines in the following format:

```
name = value
```

Comments that begin with the # character are ignored.

IBM Spectrum Protect Plus reads the template pfile and copies the entries to the new pfile that is used to start the recovered database. However, the following parameters in the template are ignored. Instead, IBM Spectrum Protect Plus sets their values to reflect appropriate values from the source database or to reflect new paths based on the renamed mount points of the recovered volumes.

- **control_files**
- **db_block_size**
- **db_create_file_dest**
- **db_recovery_file_dest**
- **log_archive_dest**
- **spfile**
- **undo_tablespace**

Additionally, cluster-related parameters like **instance_number**, **thread**, and **cluster_database** are set automatically by IBM Spectrum Protect Plus, depending on the appropriate values for the destination.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace an existing database with a database of the same name during recovery. When an Instant Disk Restore is performed for a database and another database with the same name is already running on the destination host or cluster, IBM Spectrum Protect Plus shuts down the existing database before starting up the recovered database. If this option is not selected, the restore job fails when IBM Spectrum Protect Plus detects a running database with the same name.

Continue with restores of other databases even if one fails

Toggle the recovery of a resource in a series if the previous resource recovery fails. If this option is not enabled, the restore job stops if the recovery of a resource fails.

Protocol Priority (Instant access only)

If more than one storage protocol is available, select the protocol to take priority in the job. The available protocols are **iSCSI** and **Fibre Channel**.

Mount Point Prefix

For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

- Optional: On the **Apply scripts** page, specify scripts that can be run before or after an operation runs at the job level. Batch and PowerShell scripts are supported on Windows operating systems, and shell scripts are supported on Linux operating systems.

Pre-Script

Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration > Script** page.

Post-Script

Select this check box to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration > Script** page.

Continue job/task on script error

Select this check box to continue running the job if the script that is associated with the job fails.

When you select this check box, if a pre-script or post-script completes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED.

If you clear this check box, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

- Take one of the following actions on the **Schedule** page:

- If you are running an on-demand job, click **Next**.
- If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.

- On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

An on-demand job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view the progress of the restore operation, expand the job. You can also

download the log file by clicking the download icon  .

A recurring job will begin at the scheduled start time when you start the schedule in the **Jobs and Operations > Schedule** page.

All running jobs are viewable in the **Jobs and Operations > Running Jobs** page.

What to do next

Oracle databases are always restored in non-multithreaded mode. If the databases that you restored were originally in multithreaded mode, after the restore operation is completed, you must manually configure credentials and switch the databases to the multithreaded mode.

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Adding an Oracle application server” on page 377](#)

When an Oracle application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Backing up and restoring SQL Server data

To protect content on a SQL Server server, first register the SQL Server instance so that IBM Spectrum Protect Plus recognizes it. Then create jobs for backup and restore operations.

System requirements

Ensure that your SQL Server environment meets the system requirements in [“Microsoft SQL Server database backup and restore requirements”](#) on page 22.

Registration and authentication

Register each SQL Server server in IBM Spectrum Protect Plus by name or IP address. When registering a SQL Server Cluster (AlwaysOn) node, register each node by name or IP address. Note that the IP addresses must be public-facing and listening on port 5985. The fully qualified domain name and virtual machine node DNS name must be resolvable and route-able from the IBM Spectrum Protect Plus appliance.

The user identity must have sufficient rights to install and start the IBM Spectrum Protect Plus Tools Service on the node, including the **Log on as a service** right. For more information about this right, see [Add the Log on as a service Right to an Account](#).

The default security policy uses the Windows NTLM protocol, and the user identity format follows the default *domain\name* format.

When you are using Windows group policy objects (GPO), the group policy object setting, **Network security: LAN Manager** authentication level must be set correctly. Set it with one of the following options:

- Not Defined
- Send NTLMv2 response only
- Send NTLMv2 response only. Refuse LM
- Send NTLMv2 response only. Refuse LM & NTLM

Privileges

On the SQL Server server, the system login credential must have public and sysadmin permissions enabled, plus permission to access cluster resources in a SQL Server AlwaysOn environment. If one user account is used for all SQL Server functions, a Windows login must be enabled for the SQL Server server, with public and sysadmin permissions enabled.

Every Microsoft SQL Server host can use a specific user account to access the resources of that particular SQL server instance.

To complete log backup operations, the SQL Server user registered with IBM Spectrum Protect Plus must have the sysadmin permission enabled to manage SQL Server agent jobs.

The Windows Task Scheduler is used to schedule log backups. Depending on the environment, users may receive the following error: `A specified logon session does not exist. It may already have been terminated.` This is because of a Network access Group Policy setting that needs to be disabled. For more information on how to disable this GPO, please see the following Microsoft Support article: [Task Scheduler Error “A specified logon session does not exist”](#)

Adding an SQL Server application server

When an SQL Server application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Procedure

Restriction: You can assign only one application server or file server per host. For example, if you register a host as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

To add an SQL Server host, complete the following steps.

1. In the navigation panel, click **Manage Protection > Databases > SQL**.
2. Click **Manage Application Servers**.
3. Click **Add Application Server**.
4. Populate the fields in the **Application Properties** pane:

Host Address

Enter the resolvable IP address or a resolvable path and machine name in the **Host Address** field.

Get SSL certificate thumbprint

Get the SSL certificate thumbprint for the Windows-based host. You must complete this step when registering servers for the first time or if the certificate on the server changes. This setting will only be visible if you set the global preference **Windows Clients Port (WinRM) used for application and file indexing** to 5986. For more information about global preferences, see [“Configuring global preferences” on page 173](#).

The HTTPS listener must be enabled on the host. You must create a self-signed certificate and then enable the HTTPS listener if it is not already enabled. For more information, see [How to configure WinRm for HTTPS](#).

When upgrading to IBM Spectrum Protect Plus 10.1.9, systems that are already registered in the previous version are set to trust on first use (TOFU) and the certificate thumbprint will automatically be added to the registration information in the catalog.

SSL certificate thumbprint

The SSL certificate thumbprint is displayed here. Confirm that the certificate thumbprint matches the thumbprint of the certificate on the host that you are adding.

Use existing user

Enable to select a previously entered user name and password for the provider.

UserID

Enter your user name for the provider. The user identity follows the default *domain\name* format if the virtual machine is attached to a domain. The format *local _administrator* is used if the user is a local administrator.

Password

Enter your password for the provider.

Maximum concurrent databases

Set the maximum number of databases to back up concurrently on the server. Server performance is impacted when backing up a large number of databases concurrently, as each database utilizes multiple threads and consumes bandwidth when copying data. Use this option to control the impact on server resources and minimize the impact on production operations.

5. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the application server to the IBM Spectrum Protect Plus database, and then catalogs the instance.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

What to do next

After you add the SQL Server application server, complete the following action:

Action	How to
Assign user permissions to the application server.	See “Creating a role” on page 463 .

Related concepts

[“Managing user access” on page 455](#)

By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Related tasks

[“Backing up SQL Server data” on page 392](#)

Use a backup job to back up SQL Server environments with snapshots.

[“Restoring SQL Server data” on page 396](#)

Use a restore job to restore a Microsoft SQL Server environment from snapshots. After you run IBM Spectrum Protect Plus Instant Disk Restore jobs, your SQL Server clones can be used immediately. IBM Spectrum Protect Plus catalogs and tracks all cloned instances.

Detecting SQL Server resources

SQL Server resources are automatically detected after the application server is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the application server was added.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > SQL**.
2. In the list of SQL Server instances, select an instance or click the link for the instance to navigate to the resource that you want. For example, if you want to run an inventory job for an individual database in the instance, click the instance link and then select a virtual machine.
3. Click **Run Inventory**.

Testing the connection to a SQL Server application server

You can test the connection to a SQL Server host. The test function verifies communication with the host and tests DNS settings between the IBM Spectrum Protect Plus virtual appliance and the host.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > SQL**.
2. Click **Manage Application Servers**.
3. In the list of hosts, click **Test** in the **Actions** menu for the host.

Backing up SQL Server data

Use a backup job to back up SQL Server environments with snapshots.

Before you begin

During the initial base backup, IBM Spectrum Protect Plus creates a vSnap LUN volume and creates an NTFS share on that iSCSI LUN. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus agent maps the LUN to the SQL Server server and mounts the NTFS volume to where the backup is completed. If log backups are enabled, IBM Spectrum Protect Plus creates a separate vSnap volume and creates a CIFS on that volume. Log backup transaction files are copied to this share according to the schedule created for log backup.

When the backup job is completed, the IBM Spectrum Protect Plus agent unmounts the share from the SQL Server server and creates a vSnap snapshot of the backup volume.

Review the following information:

- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,”](#) on page 455.
- Microsoft iSCSI Initiator must be enabled and running on the Windows server. An iSCSI route must be enabled between the SQL system and vSnap server. For more information, see [Microsoft iSCSI Initiator Step-by-Step Guide](#).
- IBM Spectrum Protect Plus does not support log backup of Simple recovery models.
- Failover of an SQL cluster instance during backup is not supported.
- If you plan to back up a large number of databases, you might have to increase the number of maximum worker threads on each associated SQL Server instance to ensure that backup jobs are completed successfully. The default value for maximum worker threads is 0. The server automatically determines the maximum worker threads value based on the number of processors available to the server. SQL Server uses the threads from this pool for network connections, database checkpoints, and queries. Additionally, a backup of each database requires one additional thread from this pool. If you have a large number of databases in a backup job, the default max worker threads might not be enough to back up all of the databases and the job will fail. For more information about increasing the maximum worker threads option, see [Configure the max worker threads Server Configuration Option](#).
- IBM Spectrum Protect Plus supports database backups and transaction log backups. The product name is populated in the `msdb.dbo.backupset` for records created by backups initiated from IBM Spectrum Protect Plus.
- SQL databases protected by transparent data encryption (TDE) require that a master key and certificate to be created prior to encryption. The master key and certificate are managed by the user outside of IBM Spectrum Protect Plus and are not protected as part of a backup job. Prior to restoring a TDE protected database, the master key and certificate must be restored to the restore destination so that the data protected by IBM Spectrum Protect Plus can be unencrypted. The password used during the initial encryption process must be used to unencrypt the data. For more information about moving SQL databases protected with TDE, see [Move a TDE Protected Database to Another SQL Server](#).
- For more information about log backups for SQL, see [“Log backups”](#) on page 395.

Note: Due to limitations with the Volume Shadow Copy Services (VSS) framework, leading spaces, trailing spaces, and unprintable characters should not be used in database names. For more information, see [Backing up a SQL Server database using a VSS backup application may fail for some databases](#)

Take the following actions:

- Register the SQL Servers that you want to back up. For more information, see [“Adding an SQL Server application server”](#) on page 390.
- Configure SLA policies. For more information, see [“Create backup policies”](#) on page 98.

- Before you set up and run SQL backup jobs, configure the Shadow Copy storage settings for the volumes where your SQL databases are located. This setting is configured one time for each volume. If new databases are added to the job, the setting must be configured for any new volumes that contain SQL databases. In Windows Explorer, right-click the source volume and select the **Shadow Copies** tab. Set the **Maximum size** to **No limit** or a reasonable size based on the source volume size and I/O activities, and then click **OK**. The shadow copy storage area must be on the same volume or another available volume during backup job.

Procedure

To define an SQL backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > SQL**.
2. Select an SQL Server instance to back up.

Use the search function to search for available instances and toggle the displayed instances through the **View** filter. The available options are **Standalone/Failover Cluster** and **Always On**.

Tip:

- When a resource is added to an SQL backup instance, the backup job details (name, version, SLA policy, and log backup policy) are displayed on the inventory screen.
- The **Log Backup Policy** column in the inventory screen indicates which SLA policy has log backup enabled.

3. Click **Select an SLA Policy** to add one or more SLA policies that meet your backup data criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

The job runs as defined by the SLA policies that you selected. To run the job manually, click **Jobs and Operations > Schedule**. Select the job and click **Actions > Start**.

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
 - To run an Ad hoc backup job for single resource with log backup sub policy enabled, on the **Run** wizard screen, choose the backup type (**Backup** or **Log Backup**) to start the backup job.
 - To run an Ad hoc backup job for all resources that are associated with that SLA policy has log backup sub policy enabled, on the **Run** wizard screen, select SLA policy from the drop-down list and choose the backup type (**Backup** or **Log Backup**) to start the backup job.
 - To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job”](#) on page 439.
5. Click **Select Options** to specify more options before you save the backup job.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

Enable Log Backup

Select this option to enable the backing up of transaction logs. These logs are used for recovery options such as point-in-time restore operations. If log backups are enabled for your backup jobs, transactions are continuously logged during the backup time. Notification is sent if any discontinuity is detected in log file backups.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. If you want to enable log backup for a SQL Server backup job, ensure that **Truncate SQL logs** option is not selected for a VM backup job that backs up the same databases. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

To enable log backup schedule creation for multiple databases on the same SQL Server instance, ensure that all databases are added to the same SLA policy. A staging area for the process of log backing up is not required.

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

Select one of the following options:

Back up database files one at a time using parallel streams Select this option to use parallel streams to back up your databases sequentially.

Back up database files in parallel using parallel streams Select this option to use parallel streams to backup your databases in parallel.

Finally, set the **Maximum Parallel Streams per Database** by selecting the maximum number of data streams to be used per database during the backing up process. This setting applies to each database in the job definition. Multiple databases can be backed up in parallel if the value of the option is set to **1**. Specifying Multiple parallel streams can improve backup speed in some cases.

6. Click **Save** to save the options for your backup jobs.

The job runs as defined by your SLA policy, or can be run manually from the **Job and Operations** window.

7. To configure additional options, click the **Policy Options** clipboard icon  that is associated with the job in the **SLA Policy Status** section. Set the following additional policy options:

Pre-scripts and post-scripts

Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported.

In the **Pre-script** or **Post-script** section, select an uploaded script and an application or script server where the script is due to run. To select an application server where the script runs, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script finishes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a nonzero return code, the post-script task status is reported as COMPLETED.

When this option is not enabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters in addition to the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backups operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

8. To save any additional options that you configured, click **Save**.
9. After the backup job is completed, the IBM Spectrum Protect Plus agent records the database backup status in **Windows Event log**. To navigate to the backup logs, click **Windows Event log > Applications and Services Logs > IBM Event log**.

What to do next

After you create the backup job definition, complete the following action:

Action	How to
Create an SQL Restore job definition.	See “Restoring SQL Server data” on page 396 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Starting jobs on demand” on page 433](#)

You can run any job on demand, even if the job is set to run on a schedule.

Log backups

Archived log files for databases contain committed transaction data. This transaction data can be used to run a rollforward recovery process as part of a restore operation. Using archive log backups enhances the recovery point objective for your data. Ensure that log backups are enabled in your backup jobs to allow rollforward recovery when you restore Microsoft SQL Server data.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. If enable log backup for a SQL Server backup job, ensure that **Truncate SQL logs** option is not selected for a VM backup job that backs up the same databases. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

You can perform log backups by using one of the following methods.

Common log backup approach

This method applies to all resources that are associated with an SLA policy. Select **Enable Log Backup** option to enable the backing up of transaction logs. When you enable log backups for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database. This backup creates a separate volume on the vSnap repository and the volume is mounted persistently on the SQL application server. The volume remains mounted on the SQL application server unless the **Enable Log Backup** option is cleared and a new backup job is run. To enable log backups, see topic enable log backup under step [“5” on page 393](#) in [“Backing up SQL Server data” on page 392](#).

Advanced SQL log backup approach

This method is only applicable to resources associated with an SLA policy that has log backup sub policy enabled. This log backup option is an add-on to the existing **Enable Log Backup** option assigned to the database (Inventory screen). To enable the log backup option on an existing SLA

policy, the user can edit the existing SLA policy and add the log backup sub policy. When the log backup sub policy is enabled for multiple resources in the same database, only the log backup operation on the first SLA policy is performed. This backup does not create a separate volume on the vSnap repository. For more information about log backups SLA policy for SQL, see step “5.d” on page 207 in “Creating an SLA policy for databases and file systems” on page 206.

Important: If the selected SLA policy has a backup sub policy enabled, the SLA based log backup overrides the scheduled log backup assigned to the database.

Review the following criteria before you set up log backup operations:

- To run log backups, the SQL Server agent user must be a local Windows administrator. This user must have sysadmin permission to manage SQL Server agent jobs. The agent uses that administrator account to enable and access log backup jobs. For each SQL Server instance, the SQL Server agent user also must be the user of the SQL Server service and the SQL Server agent service account. This rule is true for every SQL Server instance to be protected.
- IBM Spectrum Protect Plus does not support log backup operations for Simple recovery models.
- Avoid configuring log backups for a single SQL database by using multiple backup jobs. Logs are truncated during log backup operations. If a single SQL database is added to multiple job definitions with log backup enabled, a log backup from one job will truncate a log before the next job backs it up. This overlap might cause point-in-time restore jobs to fail.
- Before the logs are copied to the vSnap repository, IBM Spectrum Protect Plus uses the backup folder that is configured for the SQL Server instance as the staging area to collect logs. The volume where this folder is located must have sufficient space to contain the transaction logs between backup jobs. The staging area can be modified by changing the backup folder configuration in SQL Server Management Studio (SSMS).
- IBM Spectrum Protect Plus supports database backups and transaction log backups. The product name is populated in the msdb.dbo.backupset for records that are created by backups that are initiated from IBM Spectrum Protect Plus.
- IBM Spectrum Protect Plus automatically truncates post log backups of databases that it backs up. If database logs are not backed up with IBM Spectrum Protect Plus, logs are not truncated and must be managed separately.
- When an SQL backup job is completed with log backups enabled, all transaction logs up to the completion of that job are purged from the SQL Server. Log purging occurs only if the SQL backup job is completed successfully. If log backups are not backed up during a rerun of the job, log purging does not occur.
- A log backup operation for a secondary SQL Server Always On database can fail with the following error:

```
Log backup for database 'DatabaseName' on a secondary replica failed because a synchronization point could not be established on the primary database.
```

If this error occurs, change the backup preference of the availability group to Primary. Logs are then backed up from the primary replica. After a successful log backup of the primary replica is successfully completed, the backup preference can be changed.

- If a source database is overwritten, all previous transaction logs up to that point are placed in a *condense* directory after the original database is restored. When the next run of the SQL Server backup job is completed, the contents of the condense folder are removed.

Restoring SQL Server data

Use a restore job to restore a Microsoft SQL Server environment from snapshots. After you run IBM Spectrum Protect Plus Instant Disk Restore jobs, your SQL Server clones can be used immediately. IBM Spectrum Protect Plus catalogs and tracks all cloned instances.

Before you begin

Complete the following prerequisites:

- Create and run an SQL backup job. For instructions, see [“Backing up SQL Server data”](#) on page 392.
- Before an IBM Spectrum Protect Plus user can restore data, the appropriate roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations by using the **Accounts** pane. For instructions, see [Chapter 16, “Managing user access,”](#) on page 455.
- If you are planning to run a point-in-time recovery, ensure that both the restore target SQLInstance service and the IBM Spectrum Protect Plus SQL Server service use the same user account.

Review the following restrictions and considerations:

- If you are planning to run a production restore operation to an SQL Server failover cluster, the root volume of the alternative file path must be eligible to host database and log files. The volume should belong to the destination SQL Server cluster server resource group, and be a dependency of the SQL Server cluster server.
- You cannot restore data to an NTFS or FAT compressed volume because of SQL Server database restrictions. For more information, see [Description of support for SQL Server databases on compressed volumes](#).
- If you are planning to restore data to an alternative location, the SQL Server destination must be running the same version of SQL Server or a later version. For more information, see [Compatibility Support](#).
- When you are restoring data to a primary instance in an SQL Always On Availability Group environment, the database is added to the target Always On database group. After the primary restore operation, the secondary database is seeded by the SQL server in environments where automatic seeding is supported (Microsoft SQL Server 2016 and later). The database is then enabled on the destination availability group. The synchronization time depends on the amount of data that is being transferred and the connection between the primary and secondary replicas.

If automatic seeding is not supported or is not enabled, a secondary restore from the restore point with the shortest Log Sequence Number (LSN) gap of the primary instance must be completed. Log backups with the latest point-in-time restore point that is created by IBM Spectrum Protect Plus must be restored if the log backup was enabled on the primary instance. The secondary database restore operation is completed in the RESTORING state and you must issue the **T-SQL** command to add the database to the target group. For more information, see [Transact-SQL Reference \(Database Engine\)](#).

- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape before the job begins. Depending on the size of the restore, this process might take several hours.
- If you are restoring SQL Server system databases such as the master, MSDB, or model, you must restore using the Instant Access method.
 1. Complete an SQL restore with the **Instant Access** option set for the system databases.
 2. Next, stop the SQL Server instance. For more information about starting and stopping SQL Server services, see [Start, stop, pause, resume, and restart SQL Server services](#).
 3. Rename all of the data files for the system databases.
 4. Copy all data files from the Instant Access location to the location of the system databases. This is typically located at the path C:\ProgramData\SPP\mnt\id. Verify that all data files are copied.
 5. Start the SQL Server instance. Verify that the master database is successfully restored.

About this task

Instant Disk Restore uses the iSCSI protocol to immediately mount LUNs without transferring data. Databases for which snapshots were taken are cataloged and instantly recoverable with no physical transfer of data.

The following restore modes are supported:

Instant access mode

In instant access mode, no further action is taken after mounting the volume. Users can complete any custom recovery by using the files in the vSnap volume. An instant access restore of an Always On database is restored to the local destination instance. Users cannot move instance access mode to production mode.

Test mode

In test mode, the agent creates a new database by using the data files directly from the vSnap volume. Users cannot move test mode to production mode.

Production mode

In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

Procedure

To define an SQL restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection > Databases > SQL**. Click **Create job**, and then select **Restore** to open the **Restore** wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Restore > SQL**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
 - a) Click a source in the list to show the databases that are available for restore operations. You can toggle the displayed sources to show either SQL Server instances in a stand-alone or cluster environment or Always On availability groups by using the **View** filter.

You can also use the search function to search for databases in the instances or availability groups.
 - b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the minus icon  next to the item.
 - c) Click **Next** to continue.
 3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot

Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time

Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring

Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:

Option	Description
	<ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane. Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.

Option	Description
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

- On the **Restore method** page, set the restore job to run in test, production, or instant access mode by default.

For test or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

Optionally, for Production and Test restores, in the **New Database Name** field, enter the new name for the restored database. The **New Database Name** field is also displayed when you choose Production restore, but this is for restoring to a new database location on the original instance. When renaming an SQL database, the naming rules for identifiers apply. For more information, see [Database Identifiers](#). When restoring with a new name, the **Global Preferences** option **Rename SQL data and log files when database is restore in production mode with new name** must be enabled. For more information, see [“Configuring global preferences” on page 173](#).

Click **Next** to continue.

After the job is created, you can run it in test, production, or instant access mode in the **Job Sessions** pane.

- On the **Set destination** page, specify where you want to restore the database and click **Next**.

Restore to original instance

Select this option to restore the database to the original instance.

Restore to primary instance

For restore operations in an SQL Always On environment, select this option to restore the database to the primary instance of the Always On Availability Group. The database is added back to the group.

Restore to alternate instance

Select this option to restore the database to a local destination that is different from the original instance, and then select the alternative location from the list of available servers.

For restore operations in an SQL Always On environment in test mode, the source availability database is restored to the selected target instance.

For restore operations in an SQL Always On environment in production mode, the restored database is added to the target availability group if the destination instance is a primary replica. If the destination instance is a secondary replica of the target availability group, the database is restored to the secondary replica and left in restoring state.

If the automatic seeding option is enabled for the destination availability group, the secondary database file paths are synchronized with the primary database. If the primary database log is not truncated, the secondary database can be added to the availability group by SQL.

7. On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options

Set the following point-in-time recovery options:

No Recovery

Set the selected database to a RESTORING state. If you are managing transaction log backups without using IBM Spectrum Protect Plus, you can manually restore log files, and add the database to an availability group, assuming that the LSN of the secondary and primary database copies meets the criteria.

Restriction: The **No Recovery** option does not support production mode restore operations to SQL Always On groups.

Recover until end of backup

Restore the selected database to the state at the time that the backup was created.

Recover until specific point in time

When log backup is enabled by using an SQL backup job definition, point-in-time restore options will be available when you create an SQL restore job definition. Select one of the following options:

- **By Time** Select this option to configure a point-in-time recovery from a specific date and time.

Important: Selecting a date and time for which no valid log backups are available results in a failed restore. To avoid a failed restore, use the **Choose a date and time to search for a log backup coverage** option to find if there are any valid backups for a specific date and time.

- Use the **Choose a date and time to search for a log backup coverage** option to find all the available backups for a specific date and time. After you select the date and time, click **Search backups**. The UI returns all the available log backups that were created during the set date and time until end of the day.
- Use the **Choose a date and time with valid backup coverage for all of the selected databases to restore** option to select a specific point in time by entering the date and time of the available backup.

- **By Transaction ID.** Select this option to configure a point-in-time recovery by transaction ID. You can find the transaction ID in the output .json file available in the logs of the incremental backup that is performed after a log backup. To find the correct output .json file containing the recovery transaction ID, go to the job log of your incremental backup and look for the description **Performing backup of databases <DB-name> on application server <servername>**. The next entry to this description contains the name of the folder which has the correct output .json file.

Standby mode

When the Standby mode option is selected, this leaves the SQL database in a read-only state. Uncommitted transactions are undone and saved into an undo file, which may subsequently be used for bringing the database online. Transactions stored in the standby file can be applied when the database is ready to be recovered.

Note: The location of a database restored using Standby mode may be reported to be in the original database location when viewing the database in SQL Management Studio. The location is the directory that is specified by the user for a Production mode restore and the C:\ProgramData\mnt\uuid_subdirectory for a Test mode restore.

In a stand-alone restore operation, IBM Spectrum Protect Plus finds the restore points that directly proceed and follow the selected point in time. During the recovery, the older data backup volume and the newer log backup volume are mounted. If the point in time is after the last backup operation, a temporary restore point is created.

When you run restore operations in an SQL Always On environment in test mode, the restored database will join the instance where the availability group resides.

When you run restore operations in an SQL Always On environment in production mode, the restored primary database is joined to the availability group. If the automatic seeding option is enabled for the destination availability group, the secondary database file paths are synchronized with the primary database. If the primary database log is not truncated, the secondary database can be added to the availability group by SQL.

Application Options

Set the application options:

Overwrite existing database

Enable the restore job to overwrite the selected database. By default, this option is not enabled.

Tip: Before you run restore operations in an SQL Always On environment by using the production mode with the **Overwrite existing database** option, ensure that the database is not present on the replicas of the target availability group. To do so, you must manually clean up the original databases (to be overwritten) from all replicas of the target availability group.

Maximum Parallel Streams per Database

Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. If the value of the option is set to 1, multiple databases can still be restored in parallel. Multiple parallel streams might improve restore speed, but high bandwidth consumption might affect overall system performance.

This option is applicable only when you restore an SQL Server database to its original location using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace an existing database with a database of the same name during recovery. When an Instant Disk Restore is performed for a database and another database with the same name is already running on the destination host or cluster, IBM Spectrum Protect Plus shuts down the existing database before starting up the recovered database. If this option is not selected, the restore job fails when IBM Spectrum Protect Plus detects a running database with the same name.

Continue with restores of other databases even if one fails

Toggle the recovery of a resource in a series if the previous resource recovery fails. If this option is not enabled, the restore job stops if the recovery of a resource fails.

Protocol Priority (Instant Access only)

If more than one storage protocol is available, select the protocol to take priority in the job. The available protocols are **iSCSI** and **Fibre Channel**.

Mount Point Prefix

For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

- Optional: On the **Apply scripts** page, specify scripts that can be run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

Pre-Script

Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration > Script** page.

Post-Script

Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration > Script** page.

Continue job/task on script error

Select this check box to continue running the job if the script that is associated with the job fails.

When you select this check box, if a pre-script or post-script completes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED.

If you clear this check box, the backup or restore operation is not attempted, and the pre-script or post-script task status is reported as FAILED.

- Take one of the following actions on the **Schedule** page:

- If you are running an on-demand job, click **Next**.
- If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.

- On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

An on-demand job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  .

A recurring job will begin at the scheduled start time when you start the schedule in the **Jobs and Operations > Schedule** page.

All running jobs are viewable in the **Jobs and Operations > Running Jobs** page.

Related concepts

[“Configuring scripts for backup and restore operations” on page 440](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Adding an SQL Server application server” on page 390](#)

When an SQL Server application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

[“Backing up SQL Server data” on page 392](#)

Use a backup job to back up SQL Server environments with snapshots.

SAP HANA

After you successfully add your IBM SAP HANA instances to IBM Spectrum Protect Plus, you can start to protect your SAP HANA data. Create service level agreements (SLA) policies to back up and maintain SAP HANA data.

Prerequisites for SAP HANA

All prerequisites for the SAP HANA application server must be met before you start protecting SAP HANA resources with IBM Spectrum Protect Plus.

System requirements

Ensure that your SAP HANA environment meets the system requirements in [“SAP HANA requirements”](#) on page 22.

Space prerequisites

Ensure that you have enough space on the SAP HANA database management system, in the volume groups for the backup operation, and on the target volumes for copying files during the restore operation. For more information about space requirements, see [“Space requirements for SAP HANA protection”](#) on page 404. When you are restoring data to an alternative location, allocate extra dedicated volumes for the copy and restore processes. The data area on the target host is defined by the configuration parameter **basepath_datavolumes** in the `global.ini` configuration file, which is in the persistence section of the target SAP HANA instance. This setup is needed to allow copying of data from the mounted vSnap to the target host.

More configuration requirements

To meet the prerequisites for SAP HANA, complete the following checks and actions.

- Ensure that the effective file size **ulimit -f** for the IBM Spectrum Protect Plus agent user is set to **unlimited**. Alternatively, set the value to a sufficiently high value to allow copying of the largest database files in your backup and restore jobs.
- Ensure that the installed sudo version is at the recommended level. For more information, see technote [2013790](#). Then, set sudo privileges as described in [“Setting sudo privileges for SAP HANA”](#) on page 405.
- Ensure that the Linux utility package `util-linux-ng` or `util-linux` package is current.
- Unicode characters in file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.
- Ensure that your SAP HANA logical volume setup does not include nested mount points.
- Each SAP HANA instance to be protected must be registered on IBM Spectrum Protect Plus. After the instances are registered, IBM Spectrum Protect Plus runs an inventory to detect SAP HANA resources. Ensure that all instances that you want to protect are detected and listed correctly.
- Database user's credentials must be provided for each SAP HANA instance. This database user must exist in the system database of the SAP HANA instance to be protected and have the following SAP HANA system privileges: BACKUP ADMIN, INIFILE ADMIN, CATALOG READ, and DATABASE RECOVERY OPERATOR.

Space requirements for SAP HANA protection

Before you start backing up SAP HANA databases, ensure you have enough free disk space on the target and source hosts, and in the vSnap repository. Extra free disk space is required on the volume groups on the source host for creating temporary Logical Volume Manager (LVM) snapshots of the logical volumes

that the SAP HANA database files are stored on. To create LVM snapshots of a protected SAP HANA database, ensure that the volume groups with SAP HANA data have sufficient free space.

LVM snapshots

LVM snapshots are point-in-time copies of LVM logical volumes. They are space-efficient snapshots with the changed data updates from the source logical volume. LVM snapshots are created in the same volume group as the source logical volume. The IBM Spectrum Protect Plus SAP HANA agent uses LVM snapshots to create a temporary, consistent point-in-time copy of the SAP HANA database.

The IBM Spectrum Protect Plus SAP HANA agent creates an LVM snapshot which is then mounted, and copied to the vSnap repository. The duration of the file copy operation depends on the size of the SAP HANA database. During file copying, the SAP HANA application remains fully online. After the file is copied, the LVM snapshots are removed by the IBM Spectrum Protect Plus SAP HANA agent in a cleanup operation.

For every LVM snapshot logical volume containing data, allow at least 10 percent of its size as free disk space in the volume group. If the volume group has enough free disk space, the IBM Spectrum Protect Plus SAP HANA agent reserves up to 25 percent of the source logical volume size for the snapshot logical volume.

Linux LVM2

When you run the SAP HANA backup operation, SAP HANA requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system for each logical volume with data for the selected database. The logical volumes are managed by LVM2 with `lv` commands.

A software-based LVM2 snapshot is taken as a new logical volume on the same volume group. The snapshot volumes are temporarily mounted on the same machine that runs the SAP HANA instance so that they can be transferred to the vSnap repository.

The LVM2 volume manager stores the snapshot of a logical volume within the same volume group. There must be enough space on the machine to store the logical volume. The logical volume grows in size as data changes on the source volume while the snapshot exists. Ensure that the volume group has sufficient free space for the required snapshots.

Setting sudo privileges for SAP HANA

To protect your data using IBM Spectrum Protect Plus, you must install the required version of the sudo program. For the SAP HANA application server, you must set up sudo in a specific way that might be different from other application servers.

Before you begin

To determine the correct version of sudo to be installed, see technote [2013790](#).

About this task

Set up a dedicated IBM Spectrum Protect Plus agent user with the required superuser privileges for sudo. This configuration enables the agent user to run commands without a password.

Procedure

1. Create an application server user by issuing the following command:

```
useradd -m <agent>
```

where `agent` specifies the name of the IBM Spectrum Protect Plus agent user.

2. Set a password for the new user by issuing the following command:

```
passwd <agent>
```

3. To enable superuser privileges for the agent user, set the `!requiretty` setting. At the end of the sudo configuration file, add the following lines:

```
Defaults:<agent> !requiretty
<agent> ALL=(ALL) NOPASSWD:ALL
```

If your sudoers file is configured to import configurations from another directory, for example `/etc/sudoers.d`, you can add the lines in the appropriate file in that directory.

Adding an SAP HANA application server

To start protecting SAP HANA resources, you must add the server that hosts your SAP HANA instances, and set `credentials` for the instances. Repeat the procedure to add all the servers that host the SAP HANA resources.

Before you begin

IBM Spectrum Protect Plus agent secures the connections that are established by the `hdbsql` client even if they are inside the SAP HANA application server. If you are not using the default Personal Security Environment (PSE) file or if you want to specify the SSL options, you must configure the additional parameters in the `/etc/guestapps.conf` file.

The shown values represent the defaults that are applied in the absence of the file.

```
[DEFAULT]
HANATolerateUnencryptedHdbsqlConnections = True
HANAHdbsqlSSLOptions = (Empty)
```

To configure the parameters, complete the following steps:

1. Create the `guestapps.conf` file in `/etc/`, which does not exist by default.
2. Change the **HANATolerateUnencryptedHdbsqlConnections** parameter value to **False**. The default value is **True**.

Note: The SAP HANA `hdbsql` commands that are issued by the IBM Spectrum Protect Plus agent will always encrypt the communication, if this parameter is set to **False**.

3. Depending on the type of certificate being used, add any of the SSL related options listed to the **HANAHdbsqlSSLOptions** parameter, which will be passed to `hdbsql`. The default value is empty.
 - `ssltrustcert`
 - `ssltruststore <file name>`
 - `sslhostnameincert <hostname>`
 - `sslkeystore <file name>`
 - `sslprovider <provider name>`
 - `sslsniname <hostname>`

For example,:

```
HANAHdbsqlSSLOptions=-ssltrustcert -ssltruststore /usr/sap/PLE/HDB000/ulmhana5/sec/sapcli.pse
```

For more information about the SSL options and their arguments, refer to [SAP HANA Documentation](#).

4. Save the file.

About this task

To add an SAP HANA application server to IBM Spectrum Protect Plus, you must have the host address of the machine.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA**.
2. In the **SAP HANA** window, click **Manage Application Servers**, and click **Add Application Server** to add the host machine.
3. In the **Application Properties** form, enter the **Host Address**.
4. Obtain the server key and verify that the key type and key fingerprint match the host. Click **Get server key**.

Get server key

The SSH server key for the Linux-based host. You must complete this step when adding servers for the first time or if the key on the server changes.

Key type

The type of key for the Linux-based host is displayed. The following key types are supported:

- RSA with a minimum key size of 2048 bits
- ECDSA
- DSA

Key fingerprint

The MD5 hash of the SSH key fingerprint is displayed. Confirm that the key fingerprint matches the key fingerprint of the host that you are adding.

5. Choose to register the host with a user or an SSH key.
 - If you choose to specify a user, either select **Use existing user** or enter **user ID** and **Password**.
 - If you are using an SSH key, select **SSH key** from the menu.
 - password

Restriction: Any user that is specified must have sudo privileges set up.

The screenshot shows the 'SAP HANA' interface for 'Manage application servers'. The main section is 'Manage Application Servers' with a 'Create job' button. Below is the 'Application Properties' form. It includes fields for 'Host Address', 'Get server key', 'Key type', and 'Key fingerprint'. There are radio buttons for 'User' (selected) and 'SSH Key'. Under 'User', there is a checkbox for 'Use existing user' and input fields for 'User ID' (containing 'domain/user') and 'Password'. At the bottom, there is a 'Get Instances' button and a table with columns 'Name', 'Status', and 'Configured'. 'Cancel' and 'Save' buttons are at the bottom left.

Figure 25. Adding an SAP HANA agent

6. Click **Get Instances** to detect and list the SAP HANA instances that are available on the host server that you are adding.

Each SAP HANA instance is listed with its connection host address, status, and an indication of whether it is configured.

7. Click **Set Credential**, and set the database user ID, and password. Alternatively, you can select to use an existing user profile.

For more information about access control, see [Managing user access](#).

When you set credentials, you assign SAP HANA user roles for the backup and restore operations with access to role-protected SAP HANA servers. The SAP HANA user that is assigned for the role-protected SAP HANA server requires the following privileges to protect resources:

- *BACKUP ADMIN*: Authorizes BACKUP and RECOVERY statements for defining and initiating backup and recovery procedures. It also authorizes changing system configuration options with respect to backup and recovery.
- *CATALOG READ*: Authorizes unfiltered access to the data in the system views that a user has already been granted the SELECT privilege on.
- *INIFILE ADMIN*: Authorizes the user to make changes to the system settings.
- *DATABASE RECOVERY OPERATOR*: Authorizes the user to copy or recover the tenant databases. It also authorizes to check whether the backups are accessible.

8. Save the form, and repeat the steps to add other SAP HANA application servers to IBM Spectrum Protect Plus.

What to do next

After you add the SAP HANA application servers to IBM Spectrum Protect Plus, an inventory is automatically run on each application server to detect the relevant databases in those instances.

To verify that the databases are added, review the job log. Go to **Jobs and Operations**. Click **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as a wildcard in the name.

Databases must be detected to ensure that they can be protected. For instructions about running a manual inventory, see [“Detecting SAP HANA resources” on page 408](#).

Detecting SAP HANA resources

After you add your SAP HANA application servers to IBM Spectrum Protect Plus, an inventory is run automatically to detect all SAP HANA instances and databases. You can run a manual inventory on any application server to detect, list, and store all SAP HANA databases for the selected host.

Before you begin

Ensure that you added your SAP HANA application servers to IBM Spectrum Protect Plus. For instructions, see [“Adding an SAP HANA application server” on page 406](#).

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA**.

Tip: To add more SAP HANA instances to the **Instances** pane, follow the instructions in [“Adding an SAP HANA application server” on page 406](#).

2. Click **Run Inventory**.

When the inventory is running, the button changes to **Inventory In Progress**. You can run an inventory on any available application servers, but you can run only one inventory process at a time.

To monitor the inventory job, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, duration, or job name. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

3. Click an instance to open a view that shows one database entry named **ALL (SYSTEMDB + x tenants)**, where SYSTEMDB represents the system database and x tenants represents the number of tenant databases that are detected for that instance.

If the number of tenant databases is incorrect, check your SAP HANA application server and rerun the inventory. In some cases, the database entry is marked as ineligible for backup; hover over the database to reveal the reason why.

Tip: To return to the list of instances, click the **Instances** link in the **Backup SAP HANA** pane.

What to do next

To start protecting SAP HANA databases that are cataloged in the selected instance, apply a service level agreement (SLA) policy to the instance. For instructions about setting an SLA policy, see [Defining an SLA policy](#).

Testing the SAP HANA connection

After you add an SAP HANA application server, you can test the connection. The test verifies communication between IBM Spectrum Protect Plus and the SAP HANA server. It also checks that the correct sudo permissions area available for the user who is running the test.

Procedure

1. In the navigation panel, click **Manage Protection > Databases > SAP HANA**.
2. In the **SAP HANA** window, click **Manage Application Servers**, and select the host address that you want to test.

A list of the SAP HANA application servers that are available is shown.

3. Click **Actions** and choose **Test** to start the verification tests for physical and remote system connections and settings.

The test report displays a list that includes tests for the physical host network configuration, and tests for the remote server installation on the host.

4. Click **OK** to close the test report. If issues are reported, fix the issues and rerun the test to verify the fixes.

Backing up SAP HANA data

You can define backup jobs to protect your SAP HANA data. To regularly back up your data, define a backup job that includes a service level agreement (SLA) policy.

Before you begin

During the initial backup operation, IBM Spectrum Protect Plus creates a vSnap volume and Network Files System (NFS) share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus SAP HANA agent mounts the share on the SAP HANA server where the backup is completed.

Review the following prerequisites before you create a backup job definition:

- Add the application servers that you want to back up. For the procedure, see [“Adding an SAP HANA application server”](#) on page 406.

- Configure an SLA Policy. For the procedure, see [“Defining a Service Level Agreement backup job” on page 411.](#)
- Before an IBM Spectrum Protect Plus user can set up backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources, and backup and restore operations, in the **Accounts** pane. For more information, see [Chapter 16, “Managing user access,” on page 455.](#)
- Make sure that you have the required privileges to protect the resources. Run the following command in SAP HANA **hdbsql** utility to grant the required privileges:

```
create user sppadmin password MyPassw0rd no force_first_password_change
grant backup admin, catalog read, inifile admin, database recovery operator to sppadmin
```

where *sppadmin* specifies the username, *MyPassw0rd* specifies the password, and *backup admin*, *catalog read*, *inifile admin*, *database recovery operator* are the required privileges.

- Point in time recovery is not supported when one or more tenant databases or SAP HANA services are added to the database instance in the period between the chosen point in time and the time that the preceding backup job ran.

Restriction: Do not run inventory jobs at the same time when backup jobs are scheduled.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA**.
2. Select the check box for the instance that you want to back up.

Under each SAP HANA instance, data to be backed up is listed as **ALL (SYSTEMDB + x tenants)** where *SYSTEMDB* represents the system database and *x* represents the number of tenant databases. Each instance in the **Instances** pane is listed by instance name, version, and the applied SLA policy.

3. Click **Select Options** to enable or disable log backup, and to specify parallel streams to minimize the time taken for large data movement in the backup operation.

Enable Log Backup

Select this check box to back up archive logs, which allows point in time restore options and recovery options. For SAP HANA log backup settings information, see [“Log backups” on page 412.](#)

Maximum Parallel Streams per Database

Enter a number of parallel streams, and then click **Save**.

Tips for configuring the option:

Review the following tips to help you configure the option for the backup job:

- To set the option for child resources to the same value as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the option value is indeterminate. If you change a value for the option, that value is used for all selected resources after you click **Save**.
- If an option value is changed from the previously saved value, the option is shown in yellow.
- To close the **Options** pane without saving changes, click **Select Options**.

The saved options are used for all backup jobs for this instance as selected.

4. To run the backup job with these options, click the instance name, select **ALL (SYSTEMDB + x tenants)** database representation, and click **Run**.

The backup job begins, and you can view the details in **Jobs and Operations > Running Jobs**.

Tip: The **Run** button is only enabled if an SLA policy is applied to the **ALL (SYSTEMDB + x tenants)** representation of the databases.

To run an on-demand backup job for multiple databases that are associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in [“Running an ad hoc backup job” on page 439.](#)

What to do next

Important: If the backup operation fails with an error, follow the procedure described in [“Troubleshooting failed backup operations for large Db2, MongoDB, and SAP HANA databases”](#) on page 473.

Defining a Service Level Agreement backup job

After your SAP HANA instances are listed, select and apply an SLA policy to start protecting your data.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA**.
2. Select the SAP HANA instance to back up all the data in that instance.
3. Click **Select an SLA Policy** to select an SLA policy, and then click **Save**.

Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy or edit an existing policy. For more information see [“Creating an SLA policy for databases and file systems”](#) on page 206.

4. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.

For more information about SLA configuration options, see [“Setting SLA configuration options”](#) on page 411.

5. To run the policy outside of the scheduled job, select the instance. Click the **Actions** button and select **Start**. The status changes to **Running** for your chosen SLA and you can follow the progress of the job in the log shown.

What to do next

After the SLA policy is saved, you can run the policy at any time by clicking **Actions** for that policy name, and selecting **Start**. The status in the log changes to show that the backup job is in the Running state.

To cancel a job that is running, click **Actions** for that policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed-up data, or **No** to discard the backup.

Note: Backups on vSnap server will expire in accordance with the SLA attributes and get deleted by the IBM Spectrum Protect Plus server's maintenance job. You do not need to inform to SAP HANA that these backups are no longer available. The SAP HANA administrator needs to ensure that the vSnap entries are maintained for the same amount of time (or longer) than they would be in the SAP HANA backup catalog's expiration settings.

Setting SLA configuration options

After you set up a service level agreement (SLA) policy for your backup job, you can choose to configure extra options for that job. Additional SLA options include running scripts, and forcing a full base backup.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job that you are configuring, click the clipboard icon  to specify additional configuration options.
If the job is already configured, click on the icon to edit the configuration.
2. Click **Pre-Script** and define the prescript configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.

3. Click **Post-Script** and define the PostScript configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.

Scripts and script servers are configured on the **System Configuration > Script** page. For more information about working with scripts, see [Configuring scripts](#).

4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.

If this option is selected, the backup or restore operation is reattempted after an initial fail, and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.

5. Skip **Exclude Resources** for SAP HANA SLA options, as you cannot specify resources to exclude. Instances are backed up rather than individual databases.
6. To create a full, new backup of an SAP HANA instance, select **Force full backup of resources**.

A full new backup of that resource is created to replace the existing backup of that resource for one occurrence only. After that the resource is backed up incrementally as before.

Log backups

Archived logs for databases contain committed transaction data. This transaction data can be used to run a database recovery when you are running a restore operation. Using archive log backups enhances the recovery point objective for your data.

Ensure that you select the **Enable Log Backups** option to allow point in time recovery when you set up a backup job or service level agreement (SLA) policy. When selected for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database. This backup creates a separate volume on the vSnap repository, which is mounted persistently on the SAP HANA application server. The backup process updates the configuration parameters **basepath_logbackup**, and **basepath_catalogbackup** in the `global.ini` configuration file, which is in the persistence section of the target SAP HANA instance. Both parameters will be updated to the NFS mount point of the vSnap volume.

In addition, the backup process also updates the following parameters:

- **log_mode** is set to **normal**
- **log_backup_using_backint** is set to **false**
- **enable_auto_log_backup** is set to **yes**

The vSnap volume is kept mounted on the SAP HANA application server unless the **Enable Log Backup** option is cleared and a new backup job is run. During a disable log backup, the mentioned parameters will be reset to their previous values.

Truncating archive log backups

IBM Spectrum Protect Plus automatically deletes older transactional logs after a successful database backup. This action ensures that the capacity of the log archive volume is not compromised by retention of older log files. These truncated log files are stored in the vSnap repository until the corresponding backup expires and is deleted. The retention of database backups is defined in the SLA policy that you select. For more information about SLA policies, see [“Defining a Service Level Agreement backup job” on page 411](#).

IBM Spectrum Protect Plus does not manage the retention of other archived log locations.

Enabling log encryption for SAP HANA data

The SAP HANA log backups cannot be protected with the IBM Spectrum Protect Plus transport encryption feature. To protect the SAP HANA log backup data, you must enable SAP HANA backup encryption on the SAP HANA application server.

Before you begin

If you are enabling the SAP HANA log backup encryption, you must read the following documentation and ensure that the requirements are met:

- The SAP HANA documentation provides information about considerations for backup encryption. For more information, see [Points to Note: SAP HANA Backup Encryption](#).

Note: If backup encryption is enabled, the database administrator must ensure that the backup encryption root keys are backed up.

- The SAP HANA documentation provides information about the prerequisites and root key handling. For more information, see [Prerequisites: Recovering an Encrypted SAP HANA Database](#).

About this task

You can enable log backup encryption on the SAP HANA applications server. To enable log backup encryption, refer to [Enable and Disable Encryption of Data and Log Backups](#). You can also enable the backup encryption by using SAP HANA Cockpit. For more information, see [Enable Encryption of Data and Log Backups using Cockpit](#).

Restoring SAP HANA data

To restore SAP HANA data from the vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. You can choose to restore data to the original instance or to an alternative instance on a different machine, and specify recovery options, and save the job.

Before you begin

Before you create a restore job for SAP HANA, ensure that the following requirements are met:

- At least one SAP HANA backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up SAP HANA data” on page 409](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,” on page 455](#).
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job begins. Depending on the size of the restore, this process could take several hours.
- The dedicated volumes with sufficient space must be allocated to the file system structure. SAP HANA must be at the same version level on the source and target hosts for all restore operations. For more information about space requirements, see [“Space requirements for SAP HANA protection” on page 404](#). For more information about prerequisites and setup, see [“Prerequisites for SAP HANA” on page 404](#).
- Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes the data area that is defined by the configuration parameter **basepath_datavolumes** in the `global.ini` configuration file, which is in the persistence section of the target SAP HANA instance.

Important: Make sure that the tenant database operating system users and groups exist on the target system and have the same user ID (UID) and group ID (GID) when restoring SAP HANA databases that are running in isolation level high. The only exception is the SYSTEMDB, which belongs to the SAP HANA instance owner. It can have a different name and group with different UID and GID on the target system.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA** and click **Create job > Restore**.

The Restore wizard opens.

2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click SAP HANA as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
3. On the **Select source** page, click an SAP HANA instance to show the databases in that instance. Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the minus icon  next to the item.

4. Click **Next** to continue.

5. In the **Source snapshot** page, choose the type of restore operation required.

- **On-Demand: Snapshot:** creates a once-off restore operation from a database snapshot. The job is not set to recur.
- **On-Demand: Point-in-Time:** creates a once-off restore operation to recover the database state that existed at a specific date and time. The job is not set to recur.
- **Recurring:** creates a recurring job that runs on a schedule and repeats.

Tip: For an **On-Demand: Snapshot**, you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job, you can select to recover until a specific point-in-time.

6. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: <ul style="list-style-type: none">• Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order:<ul style="list-style-type: none">Backup Restores data that is backed up to a vSnap server.Replication Restores data that is replicated to a vSnap server.Object Storage Restores data that is copied to a cloud service or to a repository server.

Option	Description
	<p>Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape).</p> <ul style="list-style-type: none"> Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
<p>Use alternate vSnap server for the restore job</p>	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
<p>Restore Location Type</p>	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p> <p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
<p>Select a location</p>	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>

Option	Description
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

7. Choose an appropriate **restore method** for the destination chosen for the restore operation. Click **Next** to continue.

- **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
- **Production:** In this mode, the SAP HANA application agent first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. The copied data is then used to start the database.
- **Test:** In this mode, the agent creates a new database by using the data files directly from the vSnap repository.



Attention: When restoring to the original location, choosing test mode restore will replace any existing production SAP HANA database in the data area, which is defined by the configuration parameter **basepath datavolumes**, with symbolic links leading to the NFS mount of the vSnap volume.

To restore data to the original instance, follow the instructions in [Restoring to the original instance](#). To restore data to an alternative instance, follow the instructions in [Restoring to an alternate instance](#).

8. Set the destination for the restore operation by choosing one of the following options. Click **Next** to continue.

- **Restore to original instance:** this option restores data to the original server and original instance.
- **Restore to alternate instance:** this option restores data to a different specified location, creating a copy of the data at that location.

If you are restoring data to an alternative location, choose an instance in the **Instance** table before you click **Next**. The alternative instance must be on a different machine; unsuitable instances are not available for selection.

9. In the **Job Options** page, select the recovery, application, and advanced options for the restore operation you are defining.

Tip: Recovery options are not available for instant access restore jobs.

- **No Recovery.** This option skips any recovery after the restore operation. The database remains offline until you decide whether you want to run the recovery operation manually.
- **Recover until end of backup.** This option recovers the selected database to its state at the time the backup was created.
- **Recover until specific point-in-time.** This option includes all the backup data up to a specific point-in-time. This option is available only if you enabled log backups in your SAP HANA backup job definition. Configure a point-in-time recovery by a specific date and time, for example, Jan 1,

2019 12:18:00 AM. IBM Spectrum Protect Plus finds the restore points directly before and after the selected point-in-time. During the recovery process, the older data backup volume and the newer log backup volume are mounted. If the point-in-time is after the last backup, a temporary restore point is created. This recovery option is not available if you selected a specific restore point from the list. This option is available only when you are running an on-demand point-in-time restore job.

10. Optional: In the **Job Options** page, select the application options for the restore operation you are defining.

Tip: Application options are not available for instant access restore jobs.

- **Overwrite existing databases:** Choose this option to replace all databases of the destination SAP HANA instance during the restore recovery process. If this option is not selected, the restore job fails either when the destination SAP HANA instance and its databases are running or when database files exist in the data area which is defined by the configuration parameter **basepath_datavolumes** during the restore operation. This option is only available for production mode restore operations.
 - **Maximum Parallel Streams per Database.** You can choose to run the restore operation of data in parallel streams. This option is useful if you are restoring a large database.
11. Optional: In the **Job Options** page, select the advanced options for the restore operation you are defining.
 - **Run cleanup immediately on job failure.** This option enables the automatic cleanup of backup data as part of a restore if recovery fails. This option is selected by default. Do not clear this option unless instructed by IBM® Support for troubleshooting purposes.
 - **Continue with restores of other selected databases even if one fails.** This option continues the restore operation if one database in the instance fails to be restored successfully. The process continues for all other databases that are being restored. If this option is not selected, then the restore job stops when the recovery of a resource fails.
 - **Mount point prefix.** For instant access restore operations, specify the prefix for the path where the mount point is to be directed.
 12. Choose script options in the **Apply Scripts** page, and click **Next** to continue.
 - Select **Pre-Script** to select an uploaded script, and an application or script server where the pre-script runs. To select an application server where the script runs, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Post-Script** to select an uploaded script and an application or script server where the post-script runs. To select an application server where the script runs, clear the **Use Script Server** check box. Go to the **System Configuration > Script** page to configure scripts and script servers.
 - Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns COMPLETED. If a postscript completes with a nonzero return code, the postscript task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a FAILED status.
 13. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.
 14. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

To restore data to the original instance, follow the instructions in [Restoring to the original instance](#). To restore data to an alternative instance, follow the instructions in [Restoring to an alternate instance](#).

Restoring SAP HANA data to the original instance

You can restore a database backup to its original instance on the original host. You can restore to the latest backup or an earlier SAP HANA database backup version. This restore option runs a full production restoration of data, and existing data is overwritten at the target site if the **Overwrite existing databases** option is selected.

Before you begin

Before you create a restore job for SAP HANA, ensure that the following requirements are met:

- At least one SAP HANA backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up SAP HANA data”](#) on page 409.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,”](#) on page 455.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA** and click **Create job > Restore**.

The Restore wizard opens.

2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click SAP HANA as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
3. On the **Select source** page, click a SAP HANA instance to show the databases in that instance. Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the minus icon  next to the item.

4. Click **Next** to continue.
5. In the **Source snapshot** page, choose the type of restore operation required.
 - **On-Demand: Snapshot:** creates a once-off restore operation from a database snapshot. The job is not set to recur.
 - **On-Demand: Point-in-Time:** creates a once-off restore operation to recover the database state that existed at a specific date and time. The job is not set to recur.
 - **Recurring:** creates a recurring job that runs on a schedule and repeats.

Tip: For an **On-Demand: Snapshot**, you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job, you can select to recover until a specific point-in-time.

- Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	<p>All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:</p> <ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <p>Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane.</p>

Option	Description
	<p>Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane.</p> <p>Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.</p>
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

7. Choose an appropriate **restore method** for the destination chosen for the restore operation. Click **Next** to continue.

- **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
- **Production:** In this mode, the SAP HANA application agent first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. The copied data is then used to start the database.
- **Test:** In this mode, the agent creates a new database by using the data files directly from the vSnap repository.



Attention: When restoring to the original location, choosing test mode restore will replace any existing production SAP HANA database in the data area, which is defined by the

configuration parameter **basepath datavolumes**, with symbolic links leading to the NFS mount of the vSnap volume.

8. Set the destination for the restore operation to **Restore to original instance** to restore data to the original server. Click **Next** to continue.
9. Choose options as described in “Restoring SAP HANA data” on page 413.
10. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.

11. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

Restoring SAP HANA databases to an alternative instance

You can restore an SAP HANA database to another SAP HANA instance on an alternative host. You can also choose to restore a database to an instance with a different name. This process creates an exact copy of the database on a different host in a different instance.

It is possible to restore the SAP HANA databases that are running in isolation level high to another SAP HANA instance on an alternative host. Additionally, you can choose to restore these databases to an instance with a different name. This process creates an exact copy of the database on a different host in a different instance. All restored databases are also running in isolation level high.

If you are restoring a resource to an alternative location, you can restore the same resource multiple times without specifying different target hosts.

Before you begin

Before you create a restore job for SAP HANA, ensure that the following requirements are met:

- At least one SAP HANA backup job is set up and running successfully. For instructions about setting up a backup job, see “[Backing up SAP HANA data](#)” on page 409.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 16, “Managing user access,”](#) on page 455.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job begins. Depending on the size of the restore, this process could take several hours.
- The dedicated volumes with sufficient space must be allocated to the file system structure. SAP HANA must be at the same version level on the source and target hosts for all restore operations. For more information about space requirements, see “[Space requirements for SAP HANA protection](#)” on page 404. For more information about prerequisites and setup, see “[Prerequisites for SAP HANA](#)” on page 404.
- Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes the data area that is defined by the configuration parameter **basepath_datavolumes** in the `global.ini` configuration file, which is in the persistence section of the target SAP HANA instance.

Important: Make sure that the tenant database operating system users and groups exist on the target system and have the same user ID (UID) and group ID (GID) when restoring SAP HANA databases that

are running in isolation level high. The only exception is the SYSTEMDB, which belongs to the SAP HANA instance owner. It can have a different name and group with different UID and GID on the target system.

Restriction: If data exists on the data area which is defined by the configuration parameter **basepath_datavolumes** to which you are restoring the database backup to, and the **Overwrite existing databases** option is not selected, the restore operation fails. When the **Overwrite existing databases** option is selected, any existing data is removed from the data area which is defined by the configuration parameter **basepath_datavolumes** on the alternative host.

Procedure

1. In the navigation panel, expand **Manage Protection > Databases > SAP HANA** and click **Create job > Restore**.

The Restore wizard opens.

2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click SAP HANA as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
 - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
3. On the **Select source** page, click a SAP HANA instance to show the databases in that instance. Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the minus icon  next to the item.

4. Click **Next** to continue.

5. In the **Source snapshot** page, choose the type of restore operation required.

- **On-Demand: Snapshot:** creates a once-off restore operation from a database snapshot. The job is not set to recur.
- **On-Demand: Point-in-Time:** creates a once-off restore operation to recover the database state that existed at a specific date and time. The job is not set to recur.
- **Recurring:** creates a recurring job that runs on a schedule and repeats.

Tip: For an **On-Demand: Snapshot**, you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job, you can select to recover until a specific point-in-time.

6. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions:

Option	Description
	<ul style="list-style-type: none"> Click the backup storage type that you want to restore from. The storage types that are shown depend on the types available in your environment and are shown in the following order: <ul style="list-style-type: none"> Backup Restores data that is backed up to a vSnap server. Replication Restores data that is replicated to a vSnap server. Object Storage Restores data that is copied to a cloud service or to a repository server. Archive Restores data that is copied to a cloud service archive or to a repository server archive (tape). Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types Backup, Replication, and Archive are shown, Backup is selected by default.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

Fields that are shown for an on-demand snapshot, multiple resources restore; or recurring restore. For point-in-time restore, only Site is available for Restore Location Type.

Option	Description
Restore Location Type	<p>Select a type of location from which to restore data:</p> <ul style="list-style-type: none"> Site The site to which snapshots were backed up. The site is defined in the System Configuration > Storage > Sites pane. Cloud service copy The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server copy The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Storage > Repository servers pane. Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Storage > Cloud storage pane. Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Storage > Repository servers pane.

Option	Description
Select a location	<p>If you are restoring data from a site, select one of the following restore locations:</p> <p>Primary The primary site from which to restore snapshots.</p> <p>Secondary The secondary site from which to restore snapshots.</p> <p>If you are restoring data from a cloud or repository server, select a server from the Select a location menu.</p>
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	<p>If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the Select alternate vSnap menu.</p> <p>When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.</p>

7. Choose a **restore method** appropriate for the destination chosen for the restore operation. Click **Next** to continue.
 - **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
 - **Production:** In this mode, the SAP HANA application agent first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. That copied data is then used to start the database.
 - **Test:** In this mode, the agent creates a new database by using the data files directly from the vSnap repository.
8. Set the destination for the restore operation to **Restore to alternate instance** to restore data to a different location, which you can select from the list of eligible locations. Click **Next** to continue.

When you are restoring to an alternative location, choose an instance in the **Instance** table before you click **Next**. Unsuitable target instances cannot be selected.
9. Choose options as described in [“Restoring SAP HANA data” on page 413](#).
10. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.
11. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

Chapter 13. Protecting IBM Spectrum Protect Plus

Protect the IBM Spectrum Protect Plus application by backing up the catalog. The catalog consists of data such as application configuration settings, SLA policies, registered resources, restore points, backup storage settings, and job information.

Backing up the IBM Spectrum Protect Plus catalog

Define a backup job to protect the IBM Spectrum Protect Plus catalog.

Before you begin

The backup location is determined by the SLA policy that you select for the job definition. To optimize backup jobs, ensure that an SLA policy is created specifically for backing up IBM Spectrum Protect Plus.

To reduce system load, ensure that other jobs are not scheduled to run during the IBM Spectrum Protect Plus backup job. To create an SLA policy, follow the instructions in [“Creating an SLA policy for databases and file systems”](#) on page 206.

An IBM Spectrum Protect Plus catalog can be restored to the same location, or an alternate IBM Spectrum Protect Plus location in disaster recovery scenarios.

Procedure

To back up the IBM Spectrum Protect Plus catalog:

1. In the navigation panel, click **Manage Protection > IBM Spectrum Protect Plus > Backup**.
2. Select an SLA policy to associate with the IBM Spectrum Protect Plus catalog backup operation.
3. Click **Save** to create the job definition.

Results

The job runs as defined by the SLA policies that you selected, or you can manually run the job by clicking **Jobs and Operations > Schedule**. Then, select the job in the **Schedule** tab and click **Actions > Start**. For instructions, see [“Start a backup job”](#) on page 104.

Restoring the IBM Spectrum Protect Plus catalog

A catalog restore job restores the IBM Spectrum Protect Plus catalog from a snapshot backup.



Attention: A catalog restore operation overwrites all data in the IBM Spectrum Protect Plus server instance that you are restoring to. All IBM Spectrum Protect Plus operations stop while the data is being restored. The user interface is not accessible, and all jobs that are running are canceled. Any snapshots that are created between the backup and restore operations are not saved.

You can restore a catalog to the IBM Spectrum Protect Plus server instance from which the catalog backup originated or to an alternative server instance. To restore the catalog, the vSnap server or cloud storage system that contains the catalog snapshot must be registered in the IBM Spectrum Protect Plus server instance that you are restoring to.

When a catalog restore job is started, a job session identifier (ID) is assigned. During the initial phase, the job can be monitored in the IBM Spectrum Protect Plus UI on the job management screen until the recovery step initiates the internal database restore. When the job enters this state, IBM Spectrum Protect Plus is no longer available.

During this phase, log information is written to the location: `/data/log/catalogprotection/managedb-catalogrestore-time.log`, where *time* is epoch time. Data that is contained in this log is related to the restore of the mongo configuration and recovery catalog. After the process is complete, the `virgo` service starts and the data is written to the `virgo` log. The IBM Spectrum Protect Plus user interface

is again accessible, but the complete loading of job logs is delayed. The job logs are not immediately visible in the IBM Spectrum Protect Plus user interface. An alert is generated after the job log recovery is complete.

Related concepts

[“Managing cloud storage” on page 138](#)

You can use cloud storage as primary backup storage for container workloads and the IBM Spectrum Protect Plus catalog, or as secondary storage from the vSnap server.

Related tasks

[“Registering a vSnap server as a backup storage provider” on page 40](#)

Any vSnap server that is deployed virtually or installed physically must be registered in IBM Spectrum Protect Plus so that it can be recognized as a backup storage provider.

Restoring the IBM Spectrum Protect Plus catalog from a vSnap server

Define a job to restore the IBM Spectrum Protect Plus catalog from a vSnap server.

Procedure

To restore the catalog from a vSnap server:

1. In the navigation panel, click **Manage Protection > IBM Spectrum Protect Plus > Restore**.
2. Click the **From vSnap** tab.
3. Click the vSnap server that contains the snapshot that you want to use to restore the catalog.
Available snapshots for the vSnap server are displayed.
4. Click **Restore** for the catalog snapshot that you want.
5. Review the information in the **Catalog Restore** dialog box, and then select one of the following restore modes:

Restore the catalog and suspend all scheduled jobs

The catalog is restored and all scheduled jobs are left in a suspended state. No scheduled jobs are started, which allows for the validation and testing of catalog entries and the creation of new jobs. Typically, this option is used in DevOps use cases.

Restore the catalog

The catalog is restored and all scheduled jobs continue to run as captured in the catalog backup. Typically, this option is used in disaster recovery.

6. Click **Restore**.

Restoring the IBM Spectrum Protect Plus catalog from a cloud storage system

Define a job to restore the IBM Spectrum Protect Plus catalog from a cloud storage system.

Procedure

To restore the catalog from a cloud storage system:

1. In the navigation panel, click **Manage Protection > IBM Spectrum Protect Plus > Restore**.
2. Click the **From Cloud Storage** tab.
3. Click the cloud storage system that contains the snapshot that you want to use to restore the catalog.
Available snapshots for the cloud storage system are displayed.
4. Click **Restore** for the catalog snapshot that you want.
5. Review the information in the **Catalog Restore** dialog box, and then select one of the following restore modes:

Restore the catalog and suspend all scheduled jobs

The catalog is restored and all scheduled jobs are left in a suspended state. No scheduled jobs are started, which allows for the validation and testing of catalog entries and the creation of new jobs. Typically, this option is used in DevOps use cases.

Restore the catalog

The catalog is restored and all scheduled jobs continue to run as captured in the catalog backup. Typically, this option is used in disaster recovery.

Expire in-place snapshots for container workloads

If the catalog is restored from an OpenShift Container Platform (OCP) cluster to an alternative OCP cluster, the in-place snapshots on the original cluster are expired after the restore operation completes. References to the snapshots are removed and you can no longer restore from an in-place snapshot that was created on the original OCP cluster.

6. Click **Restore**.

Managing IBM Spectrum Protect Plus restore points

You can use the **Restore Point Retention** pane to search for restore points in the IBM Spectrum Protect Plus catalog by backup job name, view their creation and expiration dates, and override the assigned retention.

Related concepts

[“Job types” on page 431](#)

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Expiring job sessions

You can expire a job session to override the snapshot retention settings that were assigned during backup creation.

About this task

Expiring a job session will not remove a snapshot and related recovery point if the snapshot is locked by a replication or copy relationship. Run the replication or copy-enabled job to change the lock to a later snapshot. The snapshot and recovery point will be removed during the next run of the maintenance job.

Procedure

To set a job session to expire:

1. In the navigation panel, click **System Configuration > Restore Points**.
2. On the Backup Sessions tab, search for the job session or restore point. Alternatively, on the Virtual Machines / Databases tab, select either Applications or Hypervisors to search for the desired catalog entry by entering the name. Names can be searched by entering partial text, using the asterisk (*) as a wildcard character, or using the question mark (?) for pattern matching.

For more information about using the search function, see [Appendix A, “Search guidelines,” on page 481](#).

3. If you are searching from the Backup Sessions tab, use filters to fine-tune your search across job types and date range when the associated backup job started.

4. Click the search icon .

5. Select the job sessions that you want to expire.

6. From the **Actions** list, select one of the following options:

- **Expire** is used to expire a single job session.
- **Expire All Job Sessions** is used to expire all unexpired job sessions for the selected job.

Note: When IBM Cloud Object Storage with a WORM policy is the destination provider type, the **Expire** and **Expire All Job Sessions** options are not listed in the **Actions** menu. In this case, IBM Spectrum Protect Plus does not control the retention and it is instead controlled by the provider.

7. To confirm the expiration, in the dialog box, click **Yes**.

Deleting resource metadata from the IBM Spectrum Protect Plus catalog

When you run an inventory job, resources are added to the IBM Spectrum Protect Plus catalog. To release space in the catalog, you can expire the metadata from the restore points that are associated with the resources.

About this task

Expiring a resource from the catalog does not remove associated snapshots from a vSnap server or secondary backup storage.

Procedure

To expire a resource from the catalog:

1. In the navigation panel, click **System Configuration > Restore Points**.
2. Click the **Virtual Machines/Databases** tab.
3. Use the filter to search by resource type, and then enter a search string to search for a resource by name.

For more information about using the search function, see [Appendix A, “Search guidelines,” on page 481](#).

4. Click the search icon .
5. Click the delete icon  that is associated with a resource.
6. To confirm the expiration, in the dialog box, click **Yes**.

Results

The catalog metadata that is associated with the resource is removed from the catalog.

Related concepts

[“Job types” on page 431](#)

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Chapter 14. Managing jobs and operations

You can manage and monitor jobs in the **Jobs and Operations** window. You can also configure scripts to run before or after jobs.

Job types

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Backup and restore jobs are user defined. After you create these jobs, you can modify the jobs at any time. Maintenance, inventory, and report jobs are predefined and not modifiable. However, you can modify the schedules of maintenance, inventory, and report jobs.

You can run all jobs on demand, even if they are set to run on a schedule. You can also hold and release jobs that are set to run on a schedule.

The following job types are available:

Backup

A backup job defines the resources that you want to back up and the service level agreement (SLA) policy or policies that you want to apply to those resources. Each SLA policy defines when the job runs. You can run the job by using the schedule that is defined by the SLA policy or you can run the job on demand.

You can also run backup jobs for a single resource or multiple selected resources that are associated with an SLA policy rather than backing up all resources that are associated with the policy.

The job name is auto generated and is constructed of the resource type followed by the SLA policy that is used for the job. For example, a backup job for SQL Server resources that are associated with the SLA policy Gold is sql_Gold.

Restore

A restore job defines the restore point that you want to restore data from. For example, if you are restoring hypervisor data, the restore point might be a virtual machine. If you are restoring application data, the restore point might be a database.

Restore jobs are ran on a schedule or on demand.

For scheduled jobs, the job name is defined by the user who creates the job.

For on-demand jobs, the job name onDemandRestore is auto generated when the job is run. On-demand restore jobs are removed from job history when the job is 3 days old and it is not in a pending or running state, and when it is not in **Active Resources**.

Maintenance

The maintenance job runs once a day to remove resources and associated objects that are created by IBM Spectrum Protect Plus when a job that is in a pending state is deleted.

The cleanup procedure reclaims space on storage devices, cleans up the IBM Spectrum Protect Plus catalog, and removes related snapshots. The maintenance job also removes cataloged data that is associated with deleted jobs.

The job name is Maintenance

Inventory

An inventory job is run automatically when you add a resource to IBM Spectrum Protect Plus. However, you can run an inventory job at any time to detect any changes that occurred since the resource was added.

The inventory job names are Default Application Server Inventory, Default Hypervisor Inventory, and Default Storage Server Inventory.

Report

A report job runs a scheduled report. The job name is the report name preceded by Report_.

Report names are similar to the following example:

```
Report_VM Backup History
```

Log backup

Log backup files for databases contain committed transaction data. This transaction data can be used to run a rollforward recovery process as part of a restore operation. Using database log backups enhances the recovery point objective for your data.

When you enable log backups for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database.

Related concepts

[“Protecting virtualized systems” on page 213](#)

You must register the virtualized systems that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the resources that are associated with the systems.

[“Protecting databases” on page 289](#)

You must register the database applications that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the databases and resources that are associated with the applications.

Related tasks

[“Creating an SLA policy for databases and file systems” on page 206](#)

You can create custom service level agreement (SLA) policies to define backup frequency, retention, replication, and copy policies that are specific for databases and file systems.

[“Running an ad hoc backup job” on page 439](#)

With an ad hoc backup job, you can select one or more resources that are associated with an SLA policy and run an on-demand backup operation for those resources.

Creating jobs and job schedules

The method for creating jobs and job schedules depends on the job type.

You can create jobs and schedules for backup and restore jobs. The following table describes the available backup and restore jobs and provides links to the steps that are required to create the jobs and job schedules or run the jobs on demand.

Maintenance jobs are created by default. Inventory and report jobs are created automatically when an inventory operation runs or when a report is scheduled.

Job type	Description	How to create the job
Backup	You can create a job definition and assign one or more service level agreement (SLA) policies to that definition. The job definition defines the resources to back up and the SLA policy defines the schedule, targets, and other options for the backup operation.	See the topics that contain instructions for backing up data by resource type in the following sections: <ul style="list-style-type: none">• Chapter 9, “Protecting virtualized systems,” on page 213• Chapter 10, “Protecting file systems,” on page 267• Chapter 11, “Protecting data on cloud systems,” on page 281• Chapter 12, “Protecting databases,” on page 289 For example, the backup topic for VMware is “Backing up VMware data” on page 220 .

Job type	Description	How to create the job
Ad hoc backup	When a job is run for the selected SLA policy, all resources that are associated with that SLA policy are included in the backup operation. If you want to back up only selected resources by using a selected SLA policy, you can run an ad hoc job, which runs the backup operation immediately.	See “Running an ad hoc backup job” on page 439.
Restore	After you have run a backup job at least once, you can run a restore job to restore the data. You can create a restore job that runs on a schedule or that runs on demand.	See the topics that contain instructions for restoring data by resource type in the following sections: <ul style="list-style-type: none"> • Chapter 9, “Protecting virtualized systems,” on page 213 • Chapter 10, “Protecting file systems,” on page 267 • Chapter 11, “Protecting data on cloud systems,” on page 281 • Chapter 12, “Protecting databases,” on page 289 For example, the restore topic for VMware is “Restoring VMware data” on page 232.

Related concepts

[“Job types”](#) on page 431

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Related tasks

[“Creating an SLA policy for databases and file systems”](#) on page 206

You can create custom service level agreement (SLA) policies to define backup frequency, retention, replication, and copy policies that are specific for databases and file systems.

Starting jobs on demand

You can run any job on demand, even if the job is set to run on a schedule.

Procedure

Complete the following steps to start a job:

1. In the navigation panel, click **Jobs and Operations**, and click the **Schedule** tab.
2. Choose the job that you want to run, click the actions menu icon , and then click **Start**.
The job is started and added to the **Running Jobs** tab.

What to do next

To view the job log for the job, select the job on the **Running Jobs** tab and click **Job Log**. To download the log for the job, click **Download.zip**.

To view all jobs that are running or ran concurrently with the job, click **Concurrent Jobs**.

Viewing jobs

View information about the jobs that are running in your environment, the job history, the active resources that are associated with restore jobs, and scheduled jobs.

About this task

Jobs are grouped on the following job pages:

Running Jobs

This page shows jobs that are running.

Job History

This page shows jobs that ran successfully, failed, or completed processing with warnings.

Active Resources

This page shows active resources that are associated with a restore job. An example of an active resource is a file system or database that is mounted as part of a restore operation.

Schedule

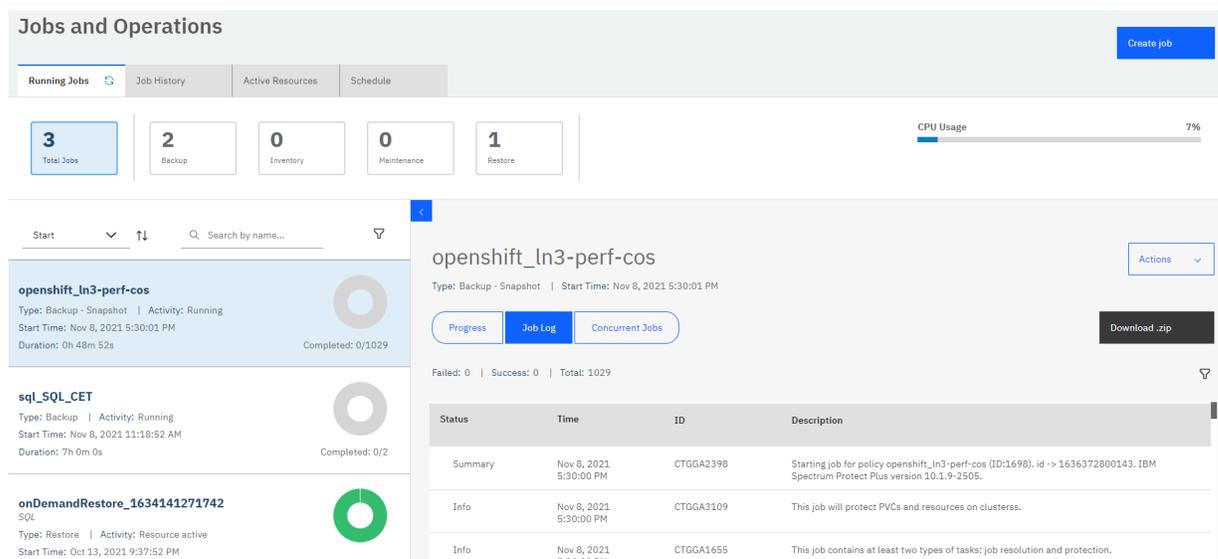
This page shows scheduled jobs.

Procedure

To view jobs, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**.
2. On the **Running Jobs** page, view the status of the jobs that are currently running, as shown in the following example.

Tip: To view the **Log Backup** job status, filter the jobs by job type by using the filter icon  and select the **Log Backup** option from the list and clear the others.



The screenshot displays the 'Jobs and Operations' interface. At the top, there are tabs for 'Running Jobs', 'Job History', 'Active Resources', and 'Schedule'. A 'Create job' button is visible in the top right. Below the tabs, there are five summary cards: 'Total Jobs' (3), 'Backup' (2), 'Inventory' (0), 'Maintenance' (0), and 'Restore' (1). A 'CPU Usage' bar shows 7% usage. The main content area shows a list of jobs. The first job is 'openshift_in3-perf-cos', which is a Backup - Snapshot job that is currently Running. It started on Nov 8, 2021 at 5:30:01 PM and has a duration of 0h 48m 52s. It is 0/1029 completed. Below this job, there are two other jobs: 'sql_SQL_CET' (Backup - Snapshot, Running, 7h 0m 0s, 0/2 completed) and 'onDemandRestore_1634141271742' (Restore - Resource active, 0/2 completed). To the right of the job list, there is a table with columns for Status, Time, ID, and Description. The table contains three rows: a Summary row for the first job, an Info row for the second job, and an Info row for the third job.

Status	Time	ID	Description
Summary	Nov 8, 2021 5:30:00 PM	CTGGA2398	Starting job for policy openshift_in3-perf-cos (ID:1698). id -> 1636372800143. IBM Spectrum Protect Plus version 10.1.9-2505.
Info	Nov 8, 2021 5:30:00 PM	CTGGA3109	This job will protect PVCs and resources on clusters.
Info	Nov 8, 2021 5:30:00 PM	CTGGA1655	This job contains at least two types of tasks: job resolution and protection.

3. To view completed jobs, click **Job History**.

Tip: To view the **Log Backup** job history, filter the jobs by job type by using the filter icon  and select the **Log Backup** option from the list and clear the others.

The ribbon across this screen shows the status of historical jobs. Use the filter to define the duration of the job history to display.

Jobs and Operations Create job

Running Jobs | **Job History** | Active Resources | Schedule

53.09% Success Rate | 81 Total Jobs | 13 Failed | 25 Warning | 43 Successful | Job history period: Last 7 days

Start | Search by name... | **sql_tapsrv04_10182**

Type: Backup | Status: Completed | Start Time: Nov 8, 2021 5:30:01 PM | Duration: 0h 10m 3s | Success: 3 | Failed: 0 | Skipped: 0 | Total: 3

Type: Backup | Status: Completed | Start Time: Nov 8, 2021 5:30:00 PM | Duration: 0h 7m 16s | Success: 1 | Failed: 0 | Skipped: 0 | Total: 1

Type: Backup | Status: Failed | Start Time: Nov 8, 2021 5:02:00 PM | Duration: 0h 23m 17s | Success: 0 | Failed: 39 | Skipped: 0 | Total: 39

sql_tapsrv04_10182 | Type: Backup | Start Time: Nov 8, 2021 5:30:01 PM | Progress | Job Log | Concurrent Jobs | Download .zip

Failed: 0 | Success: 3 | Total: 3

Status	Time	ID	Description
Summary	Nov 8, 2021 5:30:00 PM	CTGGA2398	Starting job for policy sql_tapsrv04_10182 (ID:1641). id -> 1636372800130. IBM Spectrum Protect Plus version 30.1.9-2505.
Info	Nov 8, 2021 5:30:00 PM	CTGGA2448	This job will protect databases on application servers.
Info	Nov 8, 2021 5:30:00 PM	CTGGA2449	This job contains at least two types of tasks: job resolution and protection.

- To view the active resources in your environment, click **Active Resources**, and then click **Databases**, **Virtualized Systems**, or **File Systems** to view the active resources by resource type.

To customize the columns that are displayed on each resource type, click the settings icon  to select the columns.

Jobs and Operations Create job

Running Jobs | Job History | **Active Resources** | Schedule

4 Total | 3 Databases | 1 Virtualized Systems | 0 File Systems

Resources	Type	Servers	Mount Points	Last Updated	
testdb6_archive_testmod1006	SQL	veguardian-ca6.storage.tucson.ibm.com	[veguardian-ca6.storage.tucson.ibm.com]c:\ProgramData\SPPI\mnt\76024158\	Oct 5, 2020 4:55:52 PM	
testdb4_off_testmode	SQL	veguardian-ca6.storage.tucson.ibm.com	[veguardian-ca6.storage.tucson.ibm.com]c:\ProgramData\SPPI\mnt\616588fb\	Sep 28, 2020 6:31:16 PM	
ca6_inst2_recur	SQL	veguardian-ca6.storage.tucson.ibm.com	[veguardian-ca6.storage.tucson.ibm.com]	Oct 7, 2020 3:16:31 AM	

Auto Refresh | Total: 3

- To view scheduled jobs, click **Schedule**. You can complete the following actions for scheduled jobs:
 - Start or pause a job by selecting the job and clicking **Start** or **Pause Schedule**.
 - Edit some recurring and maintenance job schedules by selecting the job and clicking the schedule icon .
 - Edit restore jobs by selecting a job and clicking click the edit icon .
 - Customize the columns that are displayed for the job table by clicking the settings icon  to select the columns.
 - Filter the jobs by job type by using the filter icon  to select the job types that you want. For example, if you want to see only backup and restore jobs, select the **Backup** and **Restore** checkboxes and clear the others.
- Optional: To download a job log and other files that reflect the information that is shown on the **Jobs and Operations** window, click **Download.zip**.

Viewing backup job progress at the resource level

View the status of individual resources in a backup job. Viewing the job at the resource level enables you to determine the backup performance of each resource. This feature provides information to help you to optimize backup performance and resolve possible issues.

About this task

This feature is available only for backup jobs. The progress of individual resources is not shown for other job types.

Procedure

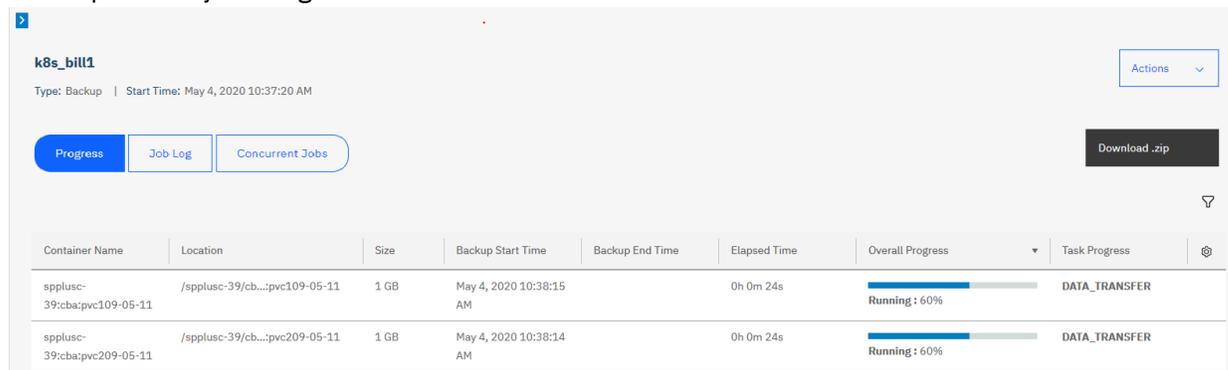
To view the progress of individual resources in a backup job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**.
2. Click **Running Jobs** for jobs that are in progress or **Job History** for jobs that are complete.
3. Select the job that contains the resources that you want to view, and then click **Progress**.

Information about each resource is shown in a table. This information includes the progress of the backup operation for each resource in the **Overall Progress** column.

If applicable for the resource type, the task that is running for the backup operation is also shown in the **Task Progress** column. This column is not included for some resource types, such as hypervisors, whose backup operations do not include individual tasks.

The following example shows the progress information for a Kubernetes backup job. In this example, the overall backup progress for the resource is 60% as shown in the **Overall Progress** column. The current backup task that is running, data transfer, is shown in the **Task Progress** column. The table was expanded by clicking the twistie .



Container Name	Location	Size	Backup Start Time	Backup End Time	Elapsed Time	Overall Progress	Task Progress
spplusc-39:cbapvc109-05-11	/spplusc-39/cb...:pvc109-05-11	1 GB	May 4, 2020 10:38:15 AM		0h 0m 24s	Running : 60%	DATA_TRANSFER
spplusc-39:cbapvc209-05-11	/spplusc-39/cb...:pvc209-05-11	1 GB	May 4, 2020 10:38:14 AM		0h 0m 24s	Running : 60%	DATA_TRANSFER

Figure 26. Viewing job information at the resource level

4. Optional: You can customize the columns that are shown in the table and filter the resources that are shown by progress status.

To customize the columns, click the settings  icon to select the columns. By default, all columns are shown.

To filter the resources by progress status, click the filter  icon and select the status values that you want. For example, if you want to see only resources that are in the process of running, select the **Running** checkbox and clear the others.

Viewing job logs

For each job run, a log is provided that shows such information as the status of the job, the start and end time for the job, and a message that is associated with the job.

Procedure

To view job logs, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**
2. Click **Running Jobs** for jobs that are in progress or **Job History** for jobs that are complete.
3. Select a job, and click **Job Log**.

The job log for the selected job is shown.

Viewing concurrent jobs

Jobs that overlap other jobs are referred to as concurrent jobs. You can view jobs that are running or ran concurrently with another job.

Procedure

To view jobs, that are running or ran concurrently with another job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**
2. Click **Running Jobs** for jobs that are in progress or **Job History** for jobs that are complete.
3. Select a job, and click **Concurrent Jobs**.

For jobs that are shown on the **Running Jobs** tab, a list of all jobs that are running concurrently with the selected job is shown. For jobs that are shown on the **Job History** tab, a list of all jobs that ran concurrently with the selected job are shown.

Restriction: Multiple backup jobs cannot back up the same resource at the same time. If multiple jobs share a resource or resources, the job that processes the resource first will run and any other jobs that start during the same time period will fail.

Pausing and resuming jobs

You can pause and resume a scheduled job. When you pause a scheduled job, the job will not run until it is resumed.

Procedure

To pause and release job schedules, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**, and click the **Schedule** tab.
2. Choose the job that you want to pause, and click the actions menu icon , and then click **Pause Schedule**.
3. To resume the job schedule, click , and then click **Release Schedule**.

Editing jobs and job schedules

You can edit the job options and schedule for some job types.

About this task

For restore jobs, you can edit the job options by using the **Restore** wizard.

For the following job types, you can edit the job schedule:

- Restore (recurring jobs)
- Inventory
- Report
- Maintenance

Procedure

To edit a job or a job schedule, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** and then click the **Schedule** tab.
2. Click the edit or schedule icon.

Option	Description
	Click this edit icon to open the Restore wizard and change the options for the job. Follow the instructions for using the wizard in the applicable resource restore topic in Chapter 9, “Protecting virtualized systems,” on page 213 and Chapter 12, “Protecting databases,” on page 289 .
	Click this edit icon to change the job schedule.

Canceling jobs

You can cancel a job that is running.

Procedure

To cancel a job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** and then click the **Running Jobs** tab.
2. Click the **Actions** menu that is associated with the job, and then click **Cancel**.

Deleting jobs

You can delete a restore or report job that has a status of IDLE.

About this task

This procedure applies only to restore and report jobs. To delete a backup job, you must delete the service level agreement (SLA) policy that is associated with that job.

Procedure

To delete a restore or report job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** and then click the **Schedule** tab.
2. Click the delete icon  that is associated with the job.

Rerunning partially completed backup jobs

If the last instance of a backup job was partially completed, you can rerun the job to back up virtual machines and databases that were skipped.

About this task

A backup job can be rerun only in the same session ID as the original partially completed backup job. No successful backup of the same resource can have completed since the partial backup job you choose to rerun.

Tip: Backup jobs can be rerun only in response to a hypervisor or database backup failure. The following events do not qualify for backup job rerun operations:

- A VM backup was completed with an FLI failure.
- A snapshot condense failure occurred for a storage system.
- A backup job failed with an unknown issue such as a cataloging error.
- A resource is missing from the vCenter.

For applications for which log backups are supported, log backups are not disabled when using the rerun feature. Log backups will be disabled for the applicable databases when the job is next started without using the on-demand backup or rerun feature.

Procedure

Complete the following steps to rerun a partially completed backup operation:

1. In the navigation panel, click **Jobs and Operations** and then click the **Job History** tab.
2. Use the search function and filters to find the last instance of the backup job that was partially completed.
3. Select the job instance and then click **Rerun**.

If the backup job cannot be rerun, the **Rerun** option is not available.

Results

All SLA options and any exclusions that are associated with the original job are included in the rerun operation. Any option or exclusion changes that you applied after the last partial backup operation are ignored. If the rerun job is completed successfully, the job summary is updated to show success.

Running an ad hoc backup job

With an ad hoc backup job, you can select one or more resources that are associated with an SLA policy and run an on-demand backup operation for those resources.

About this task

This feature associates the selected SLA policy and resources in an ad hoc job for the purposes of running an immediate, on-demand backup operation. It does not change SLA policy assignments for resources that are associated with scheduled jobs.

You can disable log backups using an ad hoc job for all application workloads that normally allow for the enablement of option. This is useful when an SLA policy has log backup enabled for several databases, but one or more specific databases must have a log backup disabled. After the backup job completes without log backup enabled, the log backup option can be re-enabled using an ad hoc backup job. For example, when attempting to clone an SQL database from an SQL Always-On Availability Group primary instance to secondary instance, the log backup on the database being cloned must be disabled prior to backing up that database. This must be done so that the cloned database in the secondary instance can be brought up in a synchronized state.

- Disable the option for log backups on the database that is to be cloned. Click on **Manage Protection > Databases**. Select the appropriate workload type.
- If necessary, change the **View** setting to make Instances visible. Select the database that is to be cloned.
- Click on **Select Options**. Deselect the **Enable Log Back up** option and then click on **Save**.
- Run an ad hoc backup using the procedure in this topic for the database.
- After the backup completes, clone the database to the secondary instance.
- Upon completion of the clone operation, re-enable the log backup option for the database using the same process.
- Finally, run a back-up job of the SLA that contains database.

Procedure

To run an ad hoc backup job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations > Create Job**.
2. Select **Ad hoc backup** to open the backup wizard.

Tips:

- You can also open the wizard from the individual hypervisor or application management pages by clicking **Manage Protection > Hypervisors** or **Manage Protection > Applications**.
 - For a running summary of your selections in the wizard, click **Preview Backup** in the navigation panel in the wizard.
3. On the **Source type** page, click the hypervisor or application for the resources that you want to include in the job.
 4. On the **Select SLA policy** page, complete the following steps:
 - a) Select **SLA policy**.
 - b) Select **Backup** or **Log Backup**.
 - c) Click **Next**.
 5. On the **Select source** page, take the following actions:
 - a) Review the available resources.

You can enter all or part of a name in the filter box to locate resources that match the search criteria. You can use the wildcard character (*) to represent all or part of a name. For example, vm2* represents all resources that begin with "vm2".
 - b) Click the plus icon  next to the resource that you want to add to the job.

To remove a resource from the list, click the minus icon  next to the resource.
 - c) Click **Next**.
 6. On the **Review** page, review the job settings and then click **Submit** to create and run the job.

What to do next

To view the status and other information about the job, click **Jobs and Operations** in the navigation panel and click the job on the **Running Jobs** tab.

Configuring scripts for backup and restore operations

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Before you begin

Review the following considerations for using scripts with hypervisors:

- The user who is running the script must have the **Log on as a service** right enabled, which is required for running prescripts and postscripts. For more information about this right, see [Add the Log on as a service Right to an Account](#).
- Windows Remote Shell (WinRM) must be enabled.

Uploading a script

Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts must be created using the associated file format for the operating system.

Procedure

Complete the following steps to upload a script:

1. In the navigation panel, click **System Configuration > Script**.
2. In the **Scripts** section, click **Upload Script**.
The **Upload Script** pane is displayed.
3. Click **Browse** to select a local script to upload.
4. Click **Save**.

The script is displayed in the **Scripts** table and can be applied to supported jobs.

What to do next

After you upload the script, complete the following action:

Action	How to
Add the script to a server from which it will run.	See “Adding a script to a server” on page 441 .

Adding a script to a server

You can add a script to the server from which the script will run.

Procedure

Complete the following steps to add a script to a server:

1. In the navigation panel, click **System Configuration > Script**.
2. In the **Script Servers** section, click **Add Script Server**.
The **Script Server Properties** pane displays.
3. Set the server options.

Host Address

Enter the resolvable IP address or a resolvable path and machine name.

Use existing user

Enable to select a previously entered user name and password for the provider.

Username

Enter your username for the provider. If entering a SQL server, the user identity follows the default *domain\name* format if the virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

Password

Enter your password for the provider.

OS Type

Select the operating system of the application server.

4. Click **Save**.

Chapter 15. Managing reports and logs

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Types of reports

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Reports are based on the data that is collected by the most recent inventory job. You can generate reports after all cataloging jobs and subsequent database condense jobs are completed. You can run the following types of reports:

- Backup storage utilization reports
- Protection reports
- System reports
- Virtual machine environment reports

Reports include interactive elements, such as searching for individual values within a report, vertical scrolling, and column sorting.

Note: You must assign tags at the VM guest level for them to be utilized for backup exclusion rules based on tags or to be used as filtering for reports in IBM Spectrum Protect Plus.

Backup storage utilization reports

IBM Spectrum Protect Plus provides backup storage utilization reports that display the storage utilization and status of your backup storage, such as vSnap servers.

To view backup storage utilization reports, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.
3. Select **Backup Storage Utilization** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** (▶) icon beside the desired report.

The following reports are available:

VM Backup Utilization report

Virtual machines can be filtered by using the **Hypervisor type**, **Hypervisor**, and **VM tags** checkboxes. The default value is **All**, which shows data for all VM backups. To specify the method to use for total size calculations, you can select the **Use storage utilization as base backup value** checkbox. When this option is selected, the total size specified in the report is calculated as the sum of the base size as reported by VMware plus any data that is transferred during incremental jobs. When this option is not selected, the total size in the report is the sum of the amount of data transferred to storage during a base backup plus any data that is transferred during incremental jobs. By default, this option is disabled.

The VM Backup Utilization report includes the VM name, its location, the hypervisor type, the SLA policy that is used to protect the VM, and the location of the backup storage used. This backup storage may be the host name or IP address of a disk, the name of the cloud server, or the name of the repository server. The backup size of each VM, and the number of recovery points that are available for each VM that is displayed. Finally, the total number of virtual machines protected appears at the bottom of the report. The **Search** box may be used to further filter report results.

vSnap Storage Utilization report

Use the report options to filter specific vSnap servers to display through the **vSnap Storage** selection box. To filter out replica destination volumes, select **Exclude Replica Destination Volumes**. For a detailed view of the individual virtual machines and databases that are protected on each vSnap server, select **Show Resources protected per vSnap Storage**. This area of the report displays the names of the virtual machines, associated hypervisor, location, and the compression and deduplication ratio of the vSnap server.

The vSnap Storage Utilization report displays the vSnap servers, the site, status, total space, free space, and used space. When expanded, the deduplication and compression ratios, if applicable, are displayed for each vSnap server. The vSnap Storage Utilization report displays both an overview of your vSnap servers and a detailed view of the individual virtual machines and databases that are protected on each vSnap server. The **Search** box may be used to further filter report results.

Note: Storage capacity and usage values that are displayed by IBM Spectrum Protect Plus might vary between those that appear on the dashboard versus those that appear on the vSnap Storage Utilization report. The dashboard displays live information, while the report reflects data from the last inventory job run. Variations are also due to differing rounding algorithms.

Related concepts

[“Report actions” on page 450](#)

You can run, create, or schedule reports in IBM Spectrum Protect Plus.

[“Types of reports” on page 443](#)

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Protection reports

IBM Spectrum Protect Plus provides reports that display the protection status of your resources. By viewing the reports and taking any necessary action, you can help to ensure that your data is protected through user-defined recovery point objective parameters.

To view protection reports, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.
3. Select **Protection** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** (▶) icon beside the selected report.

The following reports are available:

Container Persistent Volume Backup History report

The Container Persistent Volume Backup History report displays the history of persistent container volume back jobs. Use the report options to filter by Persistent Volume Claim (PVC) type and to select specific **PVCs** to display. The report can be further filtered by failed jobs or successful jobs in the **Status** field and by specific service level agreement (SLA) policies using the **SLA Policy** field. Set an integer value in the **Backup History for Past Number of Days** field to show the backup history for a specified number of days.

Database Backup History report

Run the Database Backup History report to review the protection history of specific databases. To run the report, at least one database must be specified in the **Databases** option. You can select multiple databases. Use the report options to filter **Status** by failed or successful jobs. The report can be further filtered by specific service level agreement (SLA) policies using the **SLA Policy** field. An integer value can be specified for the **Backup History for Past Number of Days** field to limit results.

In the detail view of the report, expand an associated job to view further job details, such as the reason why a job failed or the size of a successful backup. The **Search** box may be used to further filter report results.

Database SLA Policy RPO Compliance report

Use the report options to filter by **Application Type** and to select a specific **Application Server** to display. The report can be further filtered by databases that are in compliance or not in compliance with the defined RPO through the **Display Databases That Are** field, or by **Protection Type**, including data that was backed up to vSnap, using replication, using object storage copy, or using archive.

The Database SLA Policy RPO Compliance report displays databases in relation to recovery point objectives as defined in SLA policies. The quick view displays a pie chart of a count of backups to vSnap that are in compliance and those that are not in compliance. The summary view displays the SLA policy, the SLA schedule, the number of backups to vSnap that are in compliance and the number that are not in compliance, and the replications that are in compliance and not in compliance. Also displayed are databases not in compliance for the protection types which includes the database names, application servers, application types, the last successful protection time, and the out of compliance reason.

File System Backup History report

Run the File System Backup History report to review the protection history of specific file systems. To run the report, at least one server must be specified in the **Server** option and one file system must be selected for the **File System** option. Use the report options to filter **Status** by failed or successful jobs. The report can be further filtered by specific service level agreement (SLA) policies using the **SLA Policy** field. The default setting for all four options is **All**. An integer value can be specified for the **Backup History for Past Number of Days** field to limit results.

Report properties displays the creation date and the account that was used to generate the report. Also included are the report filters used when the report was generated. In the detail view of the report, the file system is listed with the server and the total number of runs. The SLA policy, time of the job, and that status of the job is displayed. The information can be expanded of an associated job to view further job details, such as the reason why a job failed and the size of a successful backup. The **Search** box may be used to further filter report results.

File System SLA Policy RPO Compliance report

Use the report options to select a specific **Server** to display. The report can be further filtered by **Protection Type**, including data that was backed up to vSnap, using replication, using object storage copy, or using archive. The default setting for these two filters is **All**. File systems that are in compliance or not in compliance with the defined RPO can be filtered through the **Display File Systems That Are** field.

The File System SLA Policy RPO Compliance report displays file systems in relation to recovery point objectives as defined in SLA policies. Report properties display the creation date and the account that was used to generate the report. Also included are the report filters used when the report was generated. The quick view displays a pie chart of a count of backups to vSnap that are in compliance and those backups that are not in compliance. The summary view displays the SLA policy, the SLA schedule, the number of backups to vSnap and jobs using replication. Non-compliant file system SLA policy jobs are included if the not in compliance filter is selected. Information that is displayed are non-compliant SLA jobs using: backup to vSnap, replication, object storage copy, and archive. For non-compliant file system SLA policy jobs, the SLA policy and SLA schedule is listed with each file system, server, the last successful protection time, and the out of compliance reason.

Microsoft 365 Backup History report

Run the Microsoft 365 Backup History report to review the protection history for specific tenants. To run the report, specify at least one tenant in the **Tenants** filter. To filter reports based on a particular jobs, specify a value in the **Status** field. Reports can be further filtered based on specific service level agreement (SLA) policies by specifying a value in the **SLA Policy** field. The default setting for all four options is **All**. To limit the report to a specified number of days, enter an integer in the **Backup History for Past Number of Days** field.

Report properties list the creation date and the account that was used to generate the report. Also included are the filters that were used to generate the report. In the detail view of the report, the file system is listed with the server and the total number of runs. The SLA policy, the job run time, and the job status are displayed. You can expand the report to view further details, such as the reason why

a job failed and the size of a successful backup. The **Search** box can be used to further filter report results.

Microsoft 365 SLA Policy RPO Compliance report

The Microsoft 365 SLA Policy RPO Compliance report displays tenants that meet or do not meet recovery point objectives (RPOs) as defined in SLA policies. Report properties display the creation date and the account that was used to generate the report. Report properties also include any filters that were applied.

Use the report options to select a specific tenant for the report. The report can be filtered by specifying a value in the **Protection Type** field. For example, you can include data that was backed up to a vSnap server, was replicated, was copied to IBM Spectrum Protect, or was archived. You can filter the report to show either policy settings.

The **Quick View** displays pie charts of the backup copies written to the vSnap server and show the levels of compliance and noncompliance. The **Summary View** displays the SLA policy, the SLA schedule, the number of backup operations to the vSnap server, and the number of replication jobs. Tenants that are not in compliance are listed with the number of applications that are affected. The last successful backup is listed with the reason why the tenant is not compliant.

Protected and Unprotected Databases report

Run the Protected and Unprotected Databases report to view the protection status of your databases. The report displays the total number of databases added to the IBM Spectrum Protect Plus inventory before backup jobs are started. Use the report options to filter by **Application Type**, **Application Server**, and **Application Server Type** to display. To exclude databases that are protected through hypervisor-based backup jobs, select **Hide Databases Protected as part of Hypervisor Backup**. To exclude unprotected databases in the report, select **Hide Unprotected Databases**.

The summary view displays an overview of your application server protection status, including the number of unprotected and protected databases, as well as the front end capacity of the protected databases. The front end capacity is the used capacity of a database. The detail view is displayed for each database type and provides further information including database names, application server, and hosting VM. The detail view also provides this information about unprotected databases in the detail view - unprotected databases section. The **Search** box may be used to further filter report results.

Protected and Unprotected File Systems report

Run the Protected and Unprotected File Systems report to view the protection status of your file systems. The report displays both the protected and unprotected file systems added to the IBM Spectrum Protect Plus inventory before backup jobs are started. Use the report options to filter by **Server**, **Operating System Type**, and **File System Type** to display. To exclude file systems that are protected through hypervisor-based backup jobs, select **Hide File Systems protected as part of Hypervisor Backup**. To exclude unprotected file systems in the report, select **Hide Unprotected File Systems**.

Report properties displays the creation date and the account that was used to generate the report. Also included are the report filters used when the report was generated. Summary view displays the protection status of registered file systems. Two detailed views are displayed, one for protected file systems and the other for unprotected file systems. Information is organized by File System, Path, File System Type, OS Type, and Server with the total number of protected and unprotected file systems displayed. The **Search** box may be used to further filter report results.

Protected and Unprotected Microsoft Office Users report

Run the Protected and Unprotected Microsoft Office Users report to view the protection status of the user accounts in your Microsoft 365 tenants. The report displays both the protected and unprotected Microsoft 365 tenants that were added to the IBM Spectrum Protect Plus inventory before the backup jobs ran. To filter a report based on tenants or applications, specify values in the **Tenants** or **Application** fields. To exclude unprotected accounts from the report, select **Hide Unprotected Accounts**.

If you decide to view all accounts, the report includes two areas, one for protected user accounts and the other for unprotected user accounts.

Within each of these sections, depending on whether you filtered results by application, there is a subsection for OneDrive and another for Outlook. For each of these applications, tenant proxies with details of each user account at that proxy site are listed.

You can use the **Search** option to locate information in the report.

Protected and Unprotected VMs report

Run the Protected and Unprotected VMs report to view the protection status of your virtual machines. The report displays the total number of virtual machines added to the IBM Spectrum Protect Plus inventory before backup jobs are started.

Use the report options to filter by **Hypervisor Type** and to select specific **Hypervisor/Accounts** to display. To exclude unprotected virtual machines in the report, select **Hide Unprotected VMs**. To exclude virtual machines that are not backed up to secondary backup storage, select **Show only the VMs with Object Storage Copy Backups**. **Tags** may also be used to filter reports.

The protected VMs displays an overview of your protected virtual machines, including the total number of VMs protected, the VM name, the hypervisor/account, type of hypervisor, location, and the managed capacity. The managed capacity is the used capacity of a virtual machine. The unprotected VMs provides the same information for virtual machines that are not protected. The **Search** box may be used to further filter report results.

VM Backup History report

Run the VM Backup History report to review the protection history of specific virtual machines. To run the report, at least one virtual machine must be specified in the **VMs** option. You can select multiple virtual machine names. Use the report options to filter **Status** by failed or successful jobs. The report can be further filtered by specific service level agreement (SLA) policies using the **SLA Policy** field. An integer may be specified for the **Backup History for Past Number of Days** field to limit results and **Tags** may also be used to filter the report.

The detail view displays the SLA policy used listed under the VM, account, and total number of runs. Information for each run can be expanded to list the backup data size. The protection time, status, and the backup storage used is also displayed. The **Search** box may be used to further filter report results.

VM SLA Policy RPO Compliance report

Use the report options to filter by **Type**, **Hypervisor/Account**, the **Protection Type** which includes data that was backed up to vSnap, using replication, using object storage copy, using archive or using snapshot, and to display virtual machines that are in compliance or not in compliance with the defined RPO through the **Display VMs That Are** field. There is also a filter for **Tags**.

The VM SLA Policy RPO Compliance report displays virtual machines in relation to recovery point objectives as defined in SLA policies. The quick view displays a pie chart of a count of backups to vSnap that are in compliance and those that are not in compliance. Also is a pie chart of snapshots that are in compliance and those that are not in compliance. A summary view displays the SLA policy used, the SLA schedule, the ratio of backups to vSnap for in compliance and not in compliance and the snapshot ratio of in compliance and not in compliance. Also displayed are VMs not in compliance view for each protection type. The **Search** box may be used to further filter report results.

Related concepts

[“Types of reports” on page 443](#)

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

System reports

IBM Spectrum Protect Plus provides system reports that display an in-depth view of the status of your configuration, including storage system information, jobs, and job status.

To view system reports, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.

3. Select **System** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** () icon beside the desired report.

The following reports are available:

Configuration report

Use the option, **Configuration Type**, to filter the configuration types to display. The Configuration report displays the configuration of the application servers, virtualized systems, backup storage for disk, object storage, and repository servers, VADP proxies, LDAP servers, and SMTP servers. Included in the report is the name of the resource, resource type (OS or application), provider, associated site, state and the TLS connection status. Not all options are displayed for each component in the Configuration report. The **Search** box may be used to further filter report results.

Job report

Use the report options to filter the job types by selecting the **Job Type** selection box and to display jobs that ran successfully over a period of time in the **Days Since Successful Run** selection box. The quick view displays a pie chart with the number of completed jobs, failed jobs, and other jobs. Summary view for jobs that have been run at least once displays the type of job, the number of jobs associated with that type, the number of runs, the number of completed jobs, failed jobs, and other jobs. The detail view for jobs run at least once includes the job, type, number of runs, and the number of completed jobs, failed jobs, and other jobs, the last successful run, and the success percentage. In all cases, other jobs are jobs that are aborted, partially run, are currently running, skipped, or stopped. In the detail view, click the plus () icon next to an associated job to view further job details such as the job ID, the average run time, last run time status, last run time, and the next scheduled run time if the job is scheduled, and the protected resources. At the end of the report is a detail view for jobs that have never run.

License report

Review the configuration of your IBM Spectrum Protect Plus environment in relation to licensed features. The following sections and fields display in this report:

Virtual Machine Protection

The **Total Number of VMs** field displays the total number of virtual machines protected through hypervisor backup jobs, plus the number of virtual machines hosting application databases protected through application backup jobs (not hypervisor backup jobs). The **Front End Capacity** field displays the used size of these virtual machines.

Physical Machine Protection

The **Total Number of Physical Servers** field displays the total number of physical application servers hosting databases that are protected through application backup jobs. The **Front End Capacity** field displays the used size of these physical application servers.

Office 365 Protection

The **Office 365 Protection** field displays the users protected through the Office 365 application backup job. The **Front End Capacity** field displays the total used size of the protected users.

Container Persistent Volume Protection

The **Container Persistent Volume Protection** field displays the protected container persistent volumes. The **Front End Capacity** field displays the used size of these protected container persistent volumes.

Backup Storage Utilization (vSnap)

The **Total Number of vSnap Servers** field displays the number of vSnap servers that are configured in IBM Spectrum Protect Plus as a backup destination. The **Target Capacity** field displays the total used capacity of the vSnap servers, excluding replica destination volumes.

Related concepts

[“Types of reports” on page 443](#)

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Running a VM environment report

You can run reports for your Virtual Machine (VM) environment in IBM Spectrum Protect Plus. Reports can help you to monitor the amount of free space on each hypervisor, the storage usage of logical unit numbers (LUNs), and the status of all VMs.

Procedure

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.
3. Select **VM Environment** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** (▶) icon beside the desired report.

The following reports are available:

VM Datastore report

Choose this to review the storage utilization of the datastores in your VM environment. Information that this report provides can be filtered using the **Hypervisor Type** and **Hypervisor**. The **Detail View Filter** controls the datastores to display in the detail view based on the percentage of space used. Use the **Show Only Orphaned Datastores** filter to view datastores that do not have any virtual machines assigned to them, or virtual machines that are in an inaccessible state. The reason for a datastore to be in an orphaned state is displayed in the **Datastore** field in the detail view.

Quick view displays a pie chart with the storage utilization of free and used space. The summary view displays the hypervisor, datastore count, the capacity and the free space. The detail view shows the datastores and displays orphaned datastores that have no VMs registered. Also displayed is the associated hypervisor, hypervisor type, datastore type, the capacity, the free space, and the percentage used. All three views contain the total datastores, total capacity, and the total free space. The **Search** box may be used to further filter report results.

VM LUNs report

Review the storage utilization of your virtual machine logical unit numbers (LUNs). Filters for this report type include **Hypervisor Type**, and **Hypervisors**. Use the **Show Only Orphaned Datastores** filter to view datastores that do not have any virtual machines assigned to them, or virtual machines that are in an inaccessible state.

In the report, the summary view displays the hypervisor, the number of LUNs associated with the hypervisor, and the capacity. In the detail view, the LUN name, LUN ID, storage vendor, hypervisor, the datastore or volume, the capacity, transport type, and the raw device mapping for each LUN is displayed. Both views display the total LUN count and total capacity. The **Search** box may be used to further filter report results.

VM Snapshot Sprawl report

This snapshot sprawl report displays the age, name, and number of snapshots that are used to protect your Hypervisor resources. The available report options to filter are by **Hypervisor Type**, **Hypervisor**, and **Tags**. Use the **Snapshot Creation Time** filter to display snapshots from specific periods of time.

The report contains a detail view which displays the snapshot name and snapshot creation time. Each snapshot appears under the associated VM, hypervisor, and hypervisor type. The total number of VMs and snapshots are displayed at the end of the view. The **Search** box may be used to further filter report results.

VM Sprawl report

Review the status of your virtual machines, including virtual machines that are powered off, powered on, or suspended. Run this report to view unused virtual machines, the date and time when they were powered off, and virtual machine templates. The available report options to filter are by **Hypervisor Type**, **Hypervisor**, **Days Since Last Powered Off**, **Dayces Since Last Suspended**, **Days Since Last Powered On**, and **Tags**.

The report contains the quick view which is a pie graph that displays the storage utilization based on the virtual machine's power state: powered off VMs, powered on VMs, templates, and suspended VMs. There are also detail views for each of the power states. The detail view - powered off VMs displays the VM name, the date and number of days since powered off, the associated hypervisor, the type of hypervisor, the provisioned space, and the datastore or volume. The total powered off VMs are displayed at the bottom of this view along with the total provisioned space. The detail view - suspended VMs contains the VM name, the date and number of days since the VM was suspended, the associated hypervisor, the type of hypervisor, the provisioned space, and the datastore of volume. The total number of suspended VMs and the total provisioned space is displayed at the bottom of the view. The detail view - templates contains the template names, associated hypervisor, the hypervisor type, the provisioned space, and the datastore or volume. The total templates and total provisioned space appears at the bottom of the view. The detail view - powered on VMs contains the VM name, the date and number of days that the VM has been powered on, the associated hypervisor, the hypervisor type, the provisioned space, and the datastore or volume. At the end of the view is the total number of powered on VMs and the total provisioned space. The **Search** box may be used to further filter report results.

VM Storage report

Review your virtual machines and associated datastores in this report. View associated datastores and provisioned space of the datastores. Use the report options to filter by **Hypervisor Type** and to select which **Hypervisor** to display.

The report contains a detail view which displays the VM name and the provisioned space. Each VM appears under the associated datastore or volume, hypervisor, and hypervisor type. The total number of datastores/volumes and VMs are displayed at the end of the view. The **Search** box may be used to further filter report results.

Related concepts

[“Types of reports” on page 443](#)

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Report actions

You can run, create, or schedule reports in IBM Spectrum Protect Plus.

Running a report

You can run IBM Spectrum Protect Plus reports with default parameters or run customized reports with custom parameters.

Before you begin

Custom roles that are assigned to users that run reports require that the appropriate permissions be set on that role so that the report can be viewed. For more information about roles, permission types, and permissions, see [“Managing roles” on page 461](#).

Procedure

To run a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.

3. Run the report by clicking the **Run Report** () icon beside the desired report.
 - To run the report with custom parameters, set the parameters in the **Run Report** window, and click **Run**. Parameters are unique to each report.
 - To run the report with default parameters, click **Run**.

What to do next

Review the report in the **Reports** pane.

Related concepts

[“Managing reports and logs” on page 443](#)

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Creating a custom report

You can modify predefined reports with custom parameters in IBM Spectrum Protect Plus and save the customized reports.

Procedure

To create a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.
3. Click the **Create Custom Report** () icon beside the desired report to be customized.
4. On the **Create Custom Report** window, select the **Parameters** tab. Enter a name for the report in the **Name** field, and enter a description for the custom report in the **Description** field. Set your customized parameters that relate to the selected report.

Note: Report names can include alphanumeric characters and the following symbols: \$-_.+!*'(). Spaces are not allowed in the report name.
5. Optionally, on the **Schedule** tab, check the **Define Schedule** box. If a schedule is to be defined, provide this information:
 - For **Repeats**, enter an integer value and select **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.
 - For **Start Time**, enter a date and time, and select the appropriate timezone. The default timezone that is displayed is based on browser settings.
 - Enter the e-mail address of the recipient that is to receive a copy of the report in the e-mail address field. At least one recipient must be added. If more addresses are required, click on the **Add a recipient** plus () icon.
6. Click the **Save Report** button.
7. To locate a custom report, click on the **Custom Reports** tab.
8. Click on the **Run Custom Report** () icon to run the report.
9. Optionally, to update a custom report, click the **Update Custom Report** () icon. To remove a custom report, click the **Remove Report** () icon.

What to do next

Run the custom report and review the report results.

Related concepts

[“Managing reports and logs” on page 443](#)

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Scheduling a report

You can schedule reports in IBM Spectrum Protect Plus to run at specific times.

Procedure

To schedule a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs > Reports**.
2. Click on the **Reports** tab.

3. Define a schedule for a report by clicking the **Schedule Report with default parameters** () icon beside the desired report.

Note: To schedule a report with non-default parameters, create a custom report. For more information, see [“Creating a custom report” on page 451](#).

4. The **Schedule Report with default parameters** window will appear.
 - For **Repeats**, enter an integer value and select **Subhourly, Hourly, Daily, Weekly, Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.
 - For **Start Time**, enter a date and time, and select the appropriate timezone. The default timezone that is displayed is based on your web-browser's settings.
 - Enter the e-mail address of the recipient that is to receive a copy of the report in the e-mail address field. At least one recipient must be added. If more addresses are required, click on the **Add a recipient** plus () icon.
5. Click the **Schedule** button.

What to do next

After the report runs, the recipient can review the report, which is delivered by email.

Related concepts

[“Managing reports and logs” on page 443](#)

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Collecting audit logs for actions

You can collect audit logs and search for actions that are completed in IBM Spectrum Protect Plus.

Procedure

To collect audit logs:

1. In the navigation panel, click **Reports and Logs > Audit Logs**.
2. Review a log of actions that were completed in IBM Spectrum Protect Plus. Information includes the users who completed the actions and descriptions of the actions.

3. To search for the actions of a specific user in IBM Spectrum Protect Plus, enter the user name in the user search field.
4. Optional: Expand the **Filters** section to further filter the displayed logs. Enter specific action descriptions and a date range in which the action was completed.
5. Click the search icon .
6. To download the audit log as a .csv file, click **Download**, and then select a location to save the file.

Related concepts

[“Managing user accounts” on page 466](#)

Before a user can log on to IBM Spectrum Protect Plus and use the available functions, a user account must be created in IBM Spectrum Protect Plus.

Chapter 16. Managing user access

By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

You can tailor IBM Spectrum Protect Plus for individual users, giving them access to the features and resources that they require.

Once resources are available to IBM Spectrum Protect Plus, they can be added to a resource group along with high-level IBM Spectrum Protect Plus items such as a hypervisor and individual screens.

Roles are then configured to define the actions that can be performed by the user associated with the resource group. These actions are then associated with one or more user accounts. The combination of one or more roles and a resource group is a permission set. User accounts may have more than one permission set applied.

Use the following sections of the **Accounts** pane to configure role-based access:

Resource Groups

A resource group defines the resources that are available to a user. Every resource that is added to IBM Spectrum Protect Plus can be included in a resource group, along with individual IBM Spectrum Protect Plus functions and screens. By defining resource groups, you can fine tune the user experience. For example, a resource group could include an individual hypervisor, with access to only backup and reporting functionality. When the resource group is associated with a role and a user, the user will see only the screens that are associated with backup and reporting for the assigned hypervisor.

Restriction: Do not assign a role-based access control (RBAC) user to more than one VMware resource group. Users that have been assigned to the Tag and Categories resource group and then are also assigned to either Hosts and Clusters or VMs and Templates will result in data not being displayed for the Hosts and Clusters view or the VMs and Templates view. Only information for Tags and Categories will be displayed when that is selected as a view when performing operations.

Roles

Roles define the actions that can be performed on the resources that are defined in a resource group. While a resource group defines the resources that will be made available to a user account, a role sets the permissions to interact with the resources defined in the resource group. For example, if a resource group is created that includes backup and restore jobs, the role determines how a user can interact with the jobs.

Permissions can be set to allow a user to create, view, and run the backup and restore jobs that are defined in a resource group, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Protect Plus features.

User accounts

A user account associates a resource group with a role. To enable a user to log in to IBM Spectrum Protect Plus and use its functions, you must first add the user as an individual user (referred to as a native user) or as part of an imported group of LDAP users, and then assign resource groups and roles to the user account. The account will have access to the resources and features that are defined in the resource group as well as the permissions to interact with the resources and features that are defined in the role.

Example: Assigning multiple permission sets to a user account

The combination of a resource group and role is known as a permission set. Multiple permission sets may be associated with a user account. You must first create the resource group and role and then make those part of a permission set. As an example, you can create a user account that only has access to certain screens, custom resource groups, custom roles, users, and a specific vCenter called *vCenter1*. In this example, we will create two permission sets and assign those sets to the user account.

Create the *ViewResourceGroup* for the screens to which the user should have access. In this example, add the User, Role, and Resource Group screens only. Next, create the *CreateResourceGroup* with the screens to which the user should have access. Again, select User, Role, and Resource Group screens only. For more information on creating a resource group, see [“Creating a resource group” on page 457](#).

Create empty roles *CreateRole* and *ViewRole*. For more information on creating a role, see [“Creating a role” on page 463](#).

Create the user account and add the two permission sets that follow to the account and set the password. For more information on creating an individual user, see [“Creating a user account for an individual user” on page 466](#) or see [“Creating a user account for an LDAP group” on page 466](#) for creating a user account for an LDAP group.

- Permission set 1 will consist of the *ViewResourceGroup* and the *ViewRole*.
- Permission set 2 will consist of the *CreateResourceGroup* and the *CreateRole*.

Enter a username and set the password for the user account. Click **Add new permission**. Expand **Permission 1** and select the *ViewRole* role and the *ViewResourceGroup* resource group. Click **Add new permission**. Expand **Permission 2** and select the *CreateRole* role and the *CreateResourceGroup* resource group. Click **Create User**.

Grant permissions to the user account to create a custom resource group and only view any resource groups that are created by the user. For information on editing resource groups, see [“Editing a resource group” on page 460](#) and for information on editing roles, see [“Editing a role” on page 465](#). Edit the following resource groups and roles:

- *CreateRole*: Select **Resource Group > Create, Edit, Delete** and click **Update Role**.
- *CreateResourceGroup*: Select **Accounts > Resource Group > All** and click **Add Resources**. Click **Update Resource Group**.
- *ViewRole*: Select **Resource Group > View** and click **Update Role**.

Grant permissions to the user account to create a custom role and only view the roles created by the user. Edit the following resource groups and roles:

- *CreateRole*: Select **Role > Create, Edit, Delete** and click **Update Role**.
- *CreateResourceGroup*: Select **Accounts > Role > All** and click **Add Resources**. Click **Update Resource Group**.
- *ViewRole*: Select **Role > View** and click **Update Role**.

Grant permission to the user account to create users and only view users created by the user. Edit the following resource groups and roles:

- *CreateRole*: Select **User > Create, Edit, Delete** and click **Update Role**.
- *CreateResourceGroup*: Select **Accounts > User > All** and click **Add Resources**. Click **Update Resource Group**.
- *ViewRole*: Select **User > View** and click **Update Role**.

Grant permission to the user account to add VMs from a specified vCenter to a resource group created by the user. Edit the following resource groups and roles:

- *ViewRole*: Select **Virtualized Systems > View** and click **Update Role**.
- *ViewResourceGroup*: Select **Virtualized System > VMware > Hosts and Clusters > vCenter1** and click **Add Resources**. Click **Update Resource Group**.

Note: In this example, *vCenter1* is the fictional name of a vCenter that has been registered in IBM Spectrum Protect Plus.

Managing user resource groups

A resource group defines the resources are made available to a user. Every resource added to IBM Spectrum Protect Plus can be included in a resource group, along with individual IBM Spectrum Protect Plus functions and screens.

Creating a resource group

Create a resource group to define the resources that are available to a user.

Before you begin

You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of those can be added as an application server to a given resource group.

Procedure

To create a resource group, complete the following steps:

1. In the navigation panel, click **Accounts > Resource Group**.
2. Click **Create Resource Group**. The **Create Resource Group** pane displays.
3. Enter a name for the resource group.
4. From the **I would like to create a resource group** menu, select one of the following options:

Option	Actions
New	<ol style="list-style-type: none">a. Select a resource type from the Choose a resource type menu.b. Select resource subtypes, and then click Add Resources. Resources are added to the Selected Resources view.
From template	<ol style="list-style-type: none">a. Select a resource group from the Which resource group would you like to use as a template? list. Resources from the selected template are added to the Selected Resources view.b. You can add resources by using the Choose a resource type list and its associated lists. <p>To view available resource types and their usage, see “Resource types ” on page 458.</p>

If you want to delete resources from the group, click the delete icon  that is associated with a resource or click **Delete All** to delete all resources.

5. When you are finished adding resources, click **Create resource group**.

Results

The resource group displays in the resource group table and can be associated with new and existing user accounts.

What to do next

After you add the resource group, complete the following action:

Action	How to
Create roles to define the actions that can be performed by the user account that is associated with the resource group. Roles are used to define permissions to interact with the resources that are defined in the resource group.	See “Creating a role” on page 463.

Resource types

Resource types are selected when resource groups are created and determine the resources that are available to a user assigned to a group.

The following resource types and subtypes are available:

<i>Table 11. Resource types and subtypes that can be selected for a resource group</i>		
Resource Type	Subtype	Description
Accounts	<ul style="list-style-type: none"> • Role • User • Identity 	Used to grant access to roles and users through the Accounts pane.
Database	<ul style="list-style-type: none"> • Db2 • Exchange Standalone/Failover Cluster • Exchange Database Availability Groups • MongoDB • SAP HANA • Oracle • SQL Standalone/Failover Cluster • SQL Always On 	Used to grant access to viewing individual application databases on an application server in IBM Spectrum Protect Plus.
Cloud	Microsoft 365	Used to grant access to cloud system resources.
Container	<ul style="list-style-type: none"> • Kubernetes • OpenShift 	Used to grant access to container resources.
File System	Windows	Used to grant access to file system resources.

Table 11. Resource types and subtypes that can be selected for a resource group (continued)

Resource Type	Subtype	Description
Server	<ul style="list-style-type: none"> • All • Db2 • Exchange • File systems • Kubernetes • OpenShift • MongoDB • SAP HANA • Microsoft 365 • Oracle • SQL 	Used to grant access to application servers in IBM Spectrum Protect Plus without access to individual databases.
Job	None	Used to grant access to Inventory, Backup, and Restore jobs. The Job resource group is mandatory for all Backup and Restore operations, including assigning SLA Policies to resources.
Report	<ul style="list-style-type: none"> • Backup Storage Utilization • Protection • System • VE Environment 	Used to grant access to report types and individual reports.
Screen	None	Used to grant or deny access to screens in the IBM Spectrum Protect Plus interface. If certain screens are not included in a resource group for a user, the user will not be able to access the functionality provided on the screen, regardless of the permissions granted to the user.
SLA Policy	None	Used to grant access to SLA Policies for Backup operations.

Table 11. Resource types and subtypes that can be selected for a resource group (continued)

Resource Type	Subtype	Description
System Configuration	Certificates	Used to grant access to TLS certificates to access cloud servers.
	Object Storage	Used to grant access to object storage that is defined as backup storage for copy operations.
	Disk	Used to grant access to vSnap backup storage servers.
	Keys	Used to grant access to the credentials required to access your resources. Identity functionality is available through the Accounts > Identities pane.
	LDAP	Used to grant access to LDAP servers for user registration.
	Logs	Used to grant access to viewing and downloading Audit and System logs.
	Repository Servers	Used to grant access to a repository server.
	Script	Used to grant access to uploaded prescripts and postscripts.
	Script Server	Used to grant access to script servers, where scripts are run during a Backup or Restore job.
	Site	Used to grant access to sites, which are assigned to vSnap backup storage servers.
	SMTP	Used to grant access to SMTP servers for job notifications.
	VADP Proxy	Used to grant access to VADP proxy servers.
Virtualized System	<ul style="list-style-type: none"> • VMware • Hyper-V • Amazon EC2 	Used to grant access to virtualized system resources.

Editing a resource group

You can edit a resource group to change the resources and features that are assigned to the group. Updated resource group settings take effect when user accounts that are associated with the resource group log in to IBM Spectrum Protect Plus.

Before you begin

Note the following considerations before editing a resource group:

- If you are signed in when the permissions or access rights for your user account are changed, you must sign out and sign in again for the updated permissions to take effect.
- You can edit any resource group that is not designated as **Cannot be modified**.

You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of those can be added as an application server to a given resource group.

Procedure

To edit a resource group, complete the following steps:

1. In the navigation panel, click **Accounts > Resource Group**.
2. Select a resource group and click the options icon ******* for the resource group. Click **Modify resources**.
3. Revise the resource group name, resources, or both.
4. Click **Update Resource Group**.

Deleting a resource group

You can delete any resource group that is not designated as **Cannot be modified**.

Procedure

To delete a resource group, complete the following steps:

1. In the navigation panel, click **Accounts > Resource Group**.
2. Select a resource group and click the options icon ******* for the resource group. Click **Delete resource group**.
3. Click **Yes**.

Managing roles

Roles define the actions that can be completed for the resources that are defined in a resource group. While a resource group defines the resources that are available to an account, a role sets the permissions to interact with the resources.

For example, if a resource group is created that includes backup and restore jobs, the role determines how a user can interact with the jobs. Permissions can be set to enable a user to create, view, and run the backup and restore jobs that are defined in a resource group, but not delete them.

Similarly, permissions can be set to create administrator accounts, enabling a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Protect Plus features.

The functionality of a role is dependent on a properly configured resource group. When selecting a predefined role or configuring a custom role, you must ensure that access to necessary IBM Spectrum Protect Plus operations, screens, and resources align with the proposed usage of the role.

About the SUPERUSER role: The SUPERUSER role provides the user with access to all IBM Spectrum Protect Plus functions. The SUPERUSER role can be assigned to only one account and that account is referred to as the superuser account. This superuser account and the SUPERUSER role are discussed in [“Managing the superuser account” on page 468](#).

The following user account roles are available:

Application Admin

Users with the Application Admin can complete the following actions:

- Register and modify application database resources that are delegated by an administrator
- Associate application databases to assigned SLA policies

- Complete backup and restore operations
- Run and schedule reports to which the user has access

Access to resources must be granted by an administrator through the **Accounts > Resource Groups** pane.

Backup Only

Users with the Backup Only role can complete the following actions:

- Create, view, and run backup operations
- View, create, and edit SLA policies to which the user has access

Access to resources, including specific backup jobs, must be granted by an administrator by clicking **Accounts > Resource Groups**.

OC_MONITOR_ROLE

The OC_MONITOR_ROLE is created when an OC_MONITOR user is created by the IBM Spectrum Protect Operations Center. This role and user are required by the Operations Center to connect to the IBM Spectrum Protect Plus environment. The OC_MONITOR_ROLE is used only by the OC_MONITOR user and provides permissions that are required to connect the Operations Center to IBM Spectrum Protect Plus. Do not edit this role.

Restore Only

Users with the Restore Only role can complete the following actions:

- Run, edit, and monitor restore operations.
- View, create, and edit SLA Policies to which the user has access.

Access to resources, including specific restore jobs, must be granted by an administrator through the **Accounts > Resource Groups** pane.

Self Service

Users with the Self Service role can monitor existing backup and restore operations that are delegated by an administrator.

Access to resources, including specific jobs, must be granted by an administrator through the **Accounts > Resource Groups** pane.

SYSADMIN

The SYSADMIN role is the administrator role. This role provides access to all resources and privileges.

Users with this role can add users and complete the following actions for all users other than the user who is assigned the SUPERUSER role:

- Modify and delete user accounts
- Change user passwords
- Assign user roles

VM Admin

Users with the VM Admin role can complete the following actions:

- Register and modify hypervisor resources to which the user has access
- Associate hypervisors to SLA policies
- Complete backup and restore operations
- Run and schedule reports to which the user has access

Access to resources must be granted by an administrator through the **Accounts > Resource Groups** pane.

Creating a role

Create roles to define the actions that can be completed by the user of an account that is associated with a resource group. Roles are used to define permissions to interact with the resources that are defined in the resource group.

Procedure

To create a user role, complete the following steps:

1. In the navigation panel, click **Accounts > Role**.
2. Click **Create Role**. The **Create Role** pane displays.
3. From the **I would like to create a role** list, select one of the following options:

Option	Actions
New	Select permissions to apply to the role. By default, none of the permissions are pre-selected.
From template	<ol style="list-style-type: none">a. Select a role from the Which role would you like to use as a template? menu. Permissions that are associated with the template role are selected by default.b. Select additional permissions to apply to the role, and delete permissions that are not required. To view available permissions and their usage, see “Permission types ” on page 463.

4. Enter a name for the role, and then click **Create Role**.

Results

The new role is displayed in the roles table and can be applied to new and existing user accounts.

Permission types

Permission types are selected when user accounts are created and determine the permissions that are available to the user.

The following permissions are available:

Name	Permissions	Description
Application	View	Used to view individual application databases on an application server in IBM Spectrum Protect Plus.
Application Server	Register, view, edit, deregister	Used to interact with application servers, such as SQL or Oracle servers, without access to individual databases.
Certificate	Create, view, edit, delete	Used to interact with TLS certificates to access cloud servers.
Object Storage	Register, view, edit, deregister	Used to interact with object storage that is defined as backup storage for copy operations.

Name	Permissions	Description
Cloud	Register, view, edit, deregister	Used to interact with cloud servers that are defined as backup storage for copy operations.
Hypervisor	Register, view, edit, deregister, options	Used to interact with hypervisor virtual machines, such as VMware or Hyper-V virtual machines.
Identity and Keys	Create, view, edit, delete	Used to interact with the credentials required to access your resources. Identity functionality is available through the Accounts > Identities pane.
LDAP	Register, view, edit, deregister	Used to interact with LDAP servers for user registration.
Log	View	Used to view Audit and System logs.
Job	Create, view, edit, run, delete	Used to interact with Inventory, Backup, and Restore jobs. Note: If the user has permission to Run a job, then they also can Hold , Release , and Perform custom restore actions for the job.
VADP Proxy	Register, view, edit, deregister	Used to interact with VADP.
Report	Create, view, edit, delete	Used to interact with reports.
Resource Group	Create, view, edit, delete	Used to interact with resource groups, which define the IBM Spectrum Protect Plus resources that are made available to a user.
Role	Create, view, edit, delete	Used to interact with roles, which define the actions that can be performed on the resources defined in a resource group.
Script	Upload, view, replace, delete	Used to interact with prescripts and postscripts that are added to IBM Spectrum Protect Plus and run before or after a job.
Script Server	Register, view, edit, deregister	Used to interact with the server on which prescripts and postscripts run.
Site	Create, view, edit, delete	Used to interact with sites, which are assigned to vSnap backup storage servers.
SMTP	Register, view, edit, deregister	Used to interact with SMTP servers for job notifications.
Backup Storage	Register, view, edit, deregister	Used to interact with vSnap backup storage servers.

Name	Permissions	Description
SLA Policy	Create, view, edit, delete	Used to interact with SLA Policies, which allow users to create customized templates for Backup jobs.
User	Create, view, edit, delete	Used to interact with users, associate a resource group with a role, and provide access to the IBM Spectrum Protect Plus user interface.

Editing a role

You can edit a role to change the resources and permissions that are assigned to the role. Updated role settings take effect when user accounts that are associated with the role log in to IBM Spectrum Protect Plus.

Before you begin

Note the following considerations before editing a role:

- If you are signed in when the permissions or access rights for your user account are changed, you must sign out and sign in again for the updated permissions to take effect.
- You can edit any role that is not designated as **Cannot be modified**.

Procedure

To edit a user role, complete the following steps

1. In the navigation panel, click **Accounts > Role**.
2. Select a role and click the options icon **☰** for the role. Click **Modify Role**.
3. Revise the role name, permissions, or both.
4. Click **Update role**.

Deleting a role

You can delete a role that is not designated as **Cannot be modified**.

Procedure

To delete a role, complete the following steps:

1. In the navigation panel, click **Accounts > Role**.
2. Select a role and click the options icon **☰** for the role. Click **Delete role**.
3. Click **Yes**.

Managing user accounts

Before a user can log on to IBM Spectrum Protect Plus and use the available functions, a user account must be created in IBM Spectrum Protect Plus.

Creating a user account for an individual user

Add an account for an individual user in IBM Spectrum Protect Plus. If you are upgrading from a version of IBM Spectrum Protect Plus that is earlier than 10.1.1, permissions assigned to users in the previous version must be reassigned in IBM Spectrum Protect Plus.

Before you begin

If you want to use custom roles and resource groups, create them before you create a user. See [“Creating a resource group”](#) on page 457 and [“Creating a role”](#) on page 463.

Procedure

To create an account for an individual user, complete the following steps:

1. In the navigation panel, click **Accounts > User**.
2. Click **Add User**. The **Add User** pane is displayed.
3. Click **Select the type of user or group you want to add > Individual new user**.
4. Enter a username and password for the user.
5. Click **Add new permission**. Expand the permission pane that appears.
6. In the **Assign Role** section, select one or more roles for the user.
7. In the **Choose Resource Groups To Assign** section, select the resource group for the user.
8. In the **Permission Groups** section, review the permissions and resources that are available to the user.
If you wish, click on **Add new permission** to add multiple permission sets to the user account.
9. Click **Create user**.

Results

The user account is displayed in the users table. Select a user from the table to view available roles, permissions, and resource groups.

Creating a user account for an LDAP group

With IBM Spectrum Protect Plus, you can use a Lightweight Directory Access Protocol (LDAP) server to manage users. When you create an LDAP user account, you can add the user account to a user group.

Before you begin

Complete the following tasks:

- Ensure that you have registered an LDAP provider with IBM Spectrum Protect Plus. To register an LDAP provider, follow the instructions in [“Adding an LDAP server”](#) on page 162.
- If you want to use custom roles and resource groups, ensure that the roles or groups are available. For instructions about creating roles and groups, see [“Creating a role”](#) on page 463 and [“Creating a resource group”](#) on page 457.

Procedure

To create a user account for an LDAP group, complete the following steps:

1. In the navigation panel, click **Accounts > User**.
2. Click **Add User**. The **Add User** pane is displayed.

3. Click **Select the type of user or group you want to add > LDAP Group**.
4. In the **Group Name** field of the **Select LDAP Group** section, specify the LDAP group by taking one of the following actions:
 - Enter the LDAP group name.
 - Search for the LDAP group name by entering partial text, an asterisk (*) as a single wildcard character, or a question mark (?) for pattern matching. To view all LDAP groups, click the **View All** button.
 - Optionally, a relative distinguished name (RDN) can be provided by filling out the **Group RDN** field.
5. LDAP Groups are displayed in **LDAP Groups** table. Select an LDAP Group.
6. Click **Add new permission**. Expand the permission pane that appears.
7. In the **Assign Role** section, select one or more roles for the user.
8. In the **Choose Resource Groups To Assign** section, select the resource group for the user.
9. In the **Permission Groups** section, review the permissions and resources that are available to the user.

If you wish, click on **Add new permission** to add multiple permission sets to the user account.
10. Click **Create user**.

Results

The user account is displayed in the users table. Optionally, to view available roles, permissions, and resource groups, select a user in the users table.

Editing a user account

You can edit the user name, password, associated resource groups, and roles for a user account, with the exception of the user who is assigned the SUPERUSER role. For information about managing the super user, see [“Managing the superuser account”](#) on page 468.

Before you begin

If you are signed in when the permissions or access rights for your user account are changed, you must sign out and sign in again for the updated permissions to take effect.

Procedure

Complete the following steps to edit the credentials of a user account:

1. In the navigation panel, click **Accounts > User**.
2. Select one or more users. If you select multiple users with different roles, you can modify only their resources and not their roles.
3. Click the options icon ******* to view available options. The options that are shown depend on the selected user or users.

Modify Permissions

Edit the associated permission sets that contain roles and a resource group associated with the user.

Modify settings

Edit the username and password, and associated permission sets.

Delete users

Delete the user account.

Change password

Change the password associated with the user account.

4. Modify the settings for the user, and then click **Update user**.

Deleting a user account

You can delete any user account, with the exception of the user who is assigned the SUPERUSER role.

Procedure

To delete a user account, complete the following steps:

1. In the navigation panel, click **Accounts > User**.
2. Select a user.
3. Click the options icon **⋮**, and then click **Delete user**.

Managing the superuser account

The superuser is the user who is assigned the IBM Spectrum Protect Plus SUPERUSER role. The SUPERUSER role provides the user with access to all IBM Spectrum Protect Plus functions.

When you log on to IBM Spectrum Protect Plus for the first time, you must log on with the username `admin` and the password `password`. You are then prompted to change this username and password. This process creates the superuser who is assigned the SUPERUSER role.

The following considerations apply to the superuser:

- There is only one superuser account.
- You cannot delete the superuser account.
- You cannot assign the SUPERUSER role to any additional user accounts that you create. You can create other accounts for administration purposes and assign the SYSADMIN role to those accounts.

Changing the superuser password

If you are logged on as the IBM Spectrum Protect Plus superuser, you can change the password for the superuser at any time. Only the superuser can change this password.

Procedure

To change the password for the superuser, complete the following steps:

1. In the navigation panel, click **Accounts > User**.
2. Select the superuser, and then click the options icon **⋮**.
3. Click **Change Password**.
4. Enter the new password in the **Password** field and current password in the **Old Password** field, and then click **Update user**.

The **Password** field is populated with the current password by default.

Resetting the superuser credentials

If you have forgotten the password for the superuser, you can reset the password. You can also reset the superuser name or maintain the existing name.

Before you begin

This procedure applies only if IBM Spectrum Protect Plus is installed as a virtual appliance. If IBM Spectrum Protect Plus is installed as a set of OpenShift containers, contact IBM Software Support for instructions about how to reset the password for the superuser.

Procedure

To reset the password for the superuser, complete the following steps:

1. Open the administrative console by entering the following URL from a supported browser:

```
https://HOSTNAME:8090/
```

where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.

2. From the **Authorization type** list, select **System**.
3. Log on as the `serveradmin` user.
4. Click **System Management**, and then click **Reset SUPERUSER password**.
5. Click **Reset Password** to confirm the request.
6. Open the IBM Spectrum Protect Plus user interface by entering the following URL from a supported browser:

```
https://host_name
```

Where *host_name* is the IP address of the virtual machine where the application is deployed.

7. Log on by entering the superuser name and the password `password`.
8. At the prompt, enter a name for the superuser and a new password. You can enter the existing name or a new name.
You cannot use the names `admin`, `root`, or `test`.
9. Click **Sign In**.

Managing identities

Some features in IBM Spectrum Protect Plus require credentials to access your resources. For example, IBM Spectrum Protect Plus connects to Oracle servers as the local operating system user that is specified during registration to complete tasks like cataloging, data protection, and data restore.

User names and passwords for your resources can be added and edited through the **Identity** pane. Then when utilizing a feature in IBM Spectrum Protect Plus that requires credentials to access a resource, select **Use existing user**, and select an identity from the drop-down menu.

Adding an identity

Add an identity to provide user credentials.

Procedure

To add an identity, complete the following steps:

1. In the navigation panel, click **Accounts > Identity**.
2. Click **Add Identity**.
3. Complete the fields in the **Identity Properties** pane:

Name

Enter a meaningful name to help identify the identity.

Username

Enter the user name that is associated with a resource, such as an SQL or Oracle server.

Password

Enter the password that is associated with a resource.

4. Click **Save**.

The identity displays in the identities table and can be selected when you are using a feature that requires credentials to access a resource through the **Use existing user** option.

Editing an identity

You can revise an identity to change the user name and password used to access an associated resource.

Procedure

To edit an identity, complete the following steps:

1. In the navigation panel, click **Accounts > Identity**.
2. Click the edit icon  that is associated with an identity.
The **Identify Properties** pane displays.
3. Revise the identity name, user name, and password.
4. Click **Save**.

The revised identity displays in the identities table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing user** option.

Deleting an identity

You can delete an identity when it becomes obsolete. If an identity is associated with a registered application server, it must be removed from the application server before it can be deleted. To remove the association, navigate to the **Backup > Manage Application Servers** page associated with the application server type, then edit the settings of the application server.

Procedure

To delete an identity, complete the following steps:

1. In the navigation panel, click **Accounts > Identity**.
2. Click the delete icon  that is associated with an identity.
3. Click **Yes** to delete the identity.

Chapter 17. Troubleshooting

Troubleshooting procedures are available for problem diagnosis and resolution.

For a list of known issues and limitations for each IBM Spectrum Protect Plus release, see [technote 567387](#).

Collecting log files for troubleshooting

To troubleshoot the IBM Spectrum Protect Plus application, you can download an archive of log files that are generated by IBM Spectrum Protect Plus.

Procedure

To collect log files for troubleshooting, complete the following steps:

1. Click the user menu, and then click **Download System Logs**.

The download process may take some time to complete.

2. Open or save the file log zip file, which contains individual log files for different IBM Spectrum Protect Plus components.

What to do next

To troubleshoot issues, complete the following steps:

1. Analyze the log files and take appropriate actions to resolve the issue.
2. If you cannot resolve the issue, submit the log files to IBM Software Support for assistance.

How do I tier data to tape or cloud storage?

You cannot tier data from IBM Spectrum Protect Plus to tape storage. You can tier data from IBM Spectrum Protect Plus to cloud storage, but only to cloud storage classes that support the rapid recall of data. When you are copying data to tape from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, it is not a good idea to use the IBM Spectrum Protect tiering function. If you are archiving data to tape, you must use a cold cache storage pool.

Review the guidelines about tape and cloud storage:

- Although you cannot tier data from IBM Spectrum Protect Plus to tape, you can archive or copy IBM Spectrum Protect Plus data to tape. To do this, define a cold-data-cache storage pool, as described in [Step 1: Creating a tape storage pool and cold-data-cache storage pool for copying data to tape](#).
- You can tier data from IBM Spectrum Protect Plus to cloud-container storage pools, but only to cloud storage classes that support the rapid recall of data. If you are using Amazon Web Services (AWS) with the Simple Storage Service (S3) protocol to move data to cloud container pools, do not move the data to Amazon S3 Glacier. For scenarios and instructions about copying or archiving data to cloud storage, see [Configuration for copying or archiving data](#). For instructions about tiering data to the cloud, see [Tiering data to cloud, tape, or file storage](#) in the IBM Spectrum Protect product documentation.

You cannot tier data from IBM Spectrum Protect Plus to tape. To store IBM Spectrum Protect Plus data on tape, copy the data to an IBM Spectrum Protect server for storage on physical tape media or in a virtual tape library. For different scenarios and more information about how to set up storage, see [“Configuration for copying or archiving data to IBM Spectrum Protect” on page 145](#) and [“Managing backup storage” on page 136](#).

To set up a cold cache storage pool for archiving or copying data to tape, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 147](#).

How does SAN work with IBM Spectrum Protect Plus and a vSnap server?

VMware production or clone restore operations can use VMware SAN transport mode, which transports data in a storage area network (SAN) environment. To run a SAN-based restore operation, you can use the advanced setting **Enable Streaming (VADP) restore**, which was introduced in IBM Spectrum Protect Plus 10.1.5. This restore operation option is set by default. Coupled with this option, you can specify SAN transport mode in the VADP proxy options for a particular site.

By using the SAN transport mode, you can restore your data by using SAN transport for the VADP transport method to read/write to the datastore over the SAN. The logical unit numbers (LUNs) that comprise that datastore must be mapped to the machine by running an initial backup. This backup operation uses the zone and LUN mask as if they were members of the vSphere cluster to access the datastore over the SAN.

Tip: To view the advanced options when you are running a production or clone restore operation, switch the job options from **Default Setup** to **Advanced Setup**.

IBM Spectrum Protect Plus restores data by creating a datastore that vSphere detects, then a storage vMotion back to the target datastore is initiated. IBM Spectrum Protect Plus does not restore data by writing directly to the datastore. For this reason, using the SAN transport mode as a communication method for block-level incremental forever processing has fewer benefits. However, for initial full backup operations, by using SAN as a transport method, works well.

For information about how to set up and run a VMware restore job, see [“Restoring VMware data” on page 232](#).

Communication

In IBM Spectrum Protect Plus, SAN backup is available through a physical proxy. Data transfer from storage to proxy is through the SAN. Communication from the proxy to the vSnap server is through the Network File System (NFS) protocol. The proxy and vSnap server can be installed on the same physical or virtual server. Review the proxy and vSnap server system requirements.

Specifying SAN as a data transport mode

To specify SAN as a transport mode, follow these steps:

1. Go to **System Configuration > VADP Proxy**. The **VADP Proxy** page opens.
2. From the table, select the server whose settings you want to edit. The **Proxy Details** pane shows the details for that server.
3. Click the actions icon  and select **Proxy Options**. The **Set VADP Proxy Options** dialog opens.

VADP Proxy

VADP Proxies

Server Name	Version	Status	Site	
spicemo...	10.1.6.409	Enabled	Secondary	
doorknob...	10.1.6.409	Unreacha...	third	
localhost	10.1.6.409	Enabled	Secondary	
cetvm79...	10.1.6.409	Enabled	Primary	

Total: 4

Auto Refresh

spicemouse12.storage.clifden.golf.ie

Proxy Details

Server Address: spicemouse12.storage.clifden.golf.ie
 Site: Secondary
 Number of Cores: 8
 Available Memory: 8.8 GB

View Tasks

Suspend
 Proxy Options
 Uninstall
 Unregister
 Edit

4. From the **Transport Modes** list, select SAN.

Tip: When selection options include multiple transport modes, the first listed mode will be used. If that mode cannot be used, the next transport mode listed for that selection option will be used for transporting the data.

5. Click **Save**.

Troubleshooting failed backup operations for large Db2, MongoDB, and SAP HANA databases

In rare cases, a backup operation for an IBM Db2 or MongoDB or SAP HANA database might fail with an error. The error can occur if the snapshot status transitions to INACTIVE for Linux Logical Volume Manager (LVM) or INVALID for AIX Enhanced Journaled File System (JFS2).

Depending on the environment, you might receive one of the following error messages.

For Db2:

```
CTGGH0209 snapshot_name is not active
```

For MongoDB:

```
CTGGI0089 snapshot_name is not active
```

For SAP HANA:

```
CTGGS0209 snapshot_name is not active
```

The instructions for troubleshooting the failed backup issue are the same for Db2, MongoDB, and SAP HANA.

To resolve the issue, complete the following tasks:

1. [“Determining the minimum acceptable snapshot size” on page 474](#)
2. [“Configuring the guestapps.conf file” on page 476](#)

Determining the minimum acceptable snapshot size

During the backup operation, if the snapshot status transitions to INACTIVE for Linux LVM or INVALID for AIX JFS2, the backup operation will end with an error.

About this task

The snapshot size must be large enough to accommodate the changes on the source logical volumes while the snapshot is being copied to the vSnap server. To fix the error, you must determine the snapshot size that doesn't result in an INACTIVE status for Linux and INVALID status for AIX by using this procedure.

If the volume group has more free space than the source volume, you can skip this procedure by setting the **Db2MaximumAllocationInPercent** parameter to 100.

Remember: When the **Db2MaximumAllocationInPercent** is set to 100, the snapshot size is adjusted to match the size of the source volume. This applies to the scenarios where all blocks on the source volume change during the backup.

Procedure

Follow the instructions for your operating system:

- To determine the minimum acceptable snapshot size for Linux, complete the following steps:
 1. Create a snapshot that is larger than 25% of the source logical volume size by issuing the following command:

```
lvcreate -s -n snapshot_name -L snapshot_size source_lv
```

where *snapshot_name* specifies the name of the snapshot, *snapshot_size* specifies the size, and *source_lv* specifies the source logical volume.

2. Periodically run the following command to monitor the status. Continue to monitor the status for a length of time that is equal to a backup operation.

```
lvs snapshot_name
```

or

```
lvdisplay snapshot_name
```

If the snapshot size is valid, you will see output that is similar to the following example:

```
lvdisplay /dev/SPPlog0vg/snap-test0
? Logical volume ?
LV Path                /dev/SPPlog0vg/snap-test0
LV Name                snap-test0
VG Name                SPPlog0vg
LV UUID                pkKCh1-C5mm-oCs8-JMfc-kTiY-F0tE-r3isw8
LV Write Access        read/write
LV Creation host, time floridaprod1, 2021-10-12 16:06:15 +0200
LV snapshot status     active destination for lvSPPlog0
LV Status               available
open                   0
LV Size                252.00 MiB
Current LE             63
COW-table size         52.00 MiB
COW-table LE          13
Allocated to snapshot  96.25%
Snapshot chunk size    4.00 KiB
Segments               1
Allocation              inherit
Read ahead sectors     auto
currently set to       8192
Block device           253:25
```

If the **LV snapshot status** field has a value of active, the snapshot size is valid. Note the snapshot size that you used for future reference.

If the snapshot size is valid, skip the remaining steps in this procedure.

If the snapshot size is invalid, you will see output that is similar to the following example. The **LV snapshot status** field shows a value of INACTIVE:

```
lvdisplay /dev/SPPllog0vg/snap-test0
? Logical volume ?
LV Path                /dev/SPPllog0vg/snap-test0
LV Name                snap-test0
VG Name                SPPllog0vg
LV UUID                pkKCh1-C5mm-oCs8-JMfc-kTiY-F0tE-r3isw8
LV Write Access        read/write
LV Creation host, time floridaprod1, 2021-10-12 16:06:15 +0200
LV snapshot status     INACTIVE destination for lvSPPllog0
LV Status              available
open                  0
LV Size                252.00 MiB
Current LE             63
COW-table size        52.00 MiB
COW-table LE          13
Snapshot chunk size   4.00 KiB
Segments              1
Allocation             inherit
Read ahead sectors    auto
currently set to      8192
Block device          253:25
```

3. If the **LV snapshot status** is INACTIVE, you must delete the snapshot by issuing the following command:

```
lvremove -f snapshot_name
```

4. Repeat steps “1” on page 474 to “3” on page 475 to determine a valid snapshot size that does not result in an INACTIVE status.
- To determine the minimum acceptable snapshot size for AIX, complete the following steps:

1. Create a snapshot that is larger than 25% of the source logical volume size by issuing the following command:

```
snapshot -o snapfrom=src_fs_path -o size=snapshot_size M
```

where *src_fs_path* specifies source file system path, and *snapshot_size* specifies the snapshot size.

2. Periodically run the following command to monitor the status. Run the command for a length of time that is equal to the time required for a backup operation.

```
snapshot -q src_fs_path
```

If the snapshot size is valid, you will see output that is similar to the following example:

```
snapshot -q /db2/SPN/log_dir/NODE0000
Snapshots for /db2/SPN/log_dir/NODE0000
Current Location      512-blocks      Free Time
* /dev/fs1v00        65536           64768 Wed Oct 13 19:41:19
CEST 2021
```

If the output does not include the word INVALID, the snapshot size is valid. Note the snapshot size that you used for future reference.

If the snapshot size is valid, skip the remaining steps in this procedure.

If the snapshot size is invalid, you will see output that is similar to the following example:

```
snapshot -q /db2/SPN/log_dir/NODE0000
Snapshots for /db2/SPN/log_dir/NODE0000
Current Location      512-blocks      Free Time
INVALID /dev/fs1v00  65536           Wed Oct 13
19:41:19 CEST 2021
```

3. If the snapshot status is INVALID, you must delete the snapshot by issuing the following command:

```
snapshot -d snapshot_LV
```

4. Repeat steps “1” on page 474 to “3” on page 475 to determine a valid snapshot size that does not result in INVALID status.

What to do next

Configure the `guestapps.conf` file by using the resulting snapshot size. For instructions, see [“Configuring the guestapps.conf file” on page 476](#).

Configuring the guestapps.conf file

To help ensure that backup operations will be successful for Db2, MongoDB, and SAP HANA databases, create the `guestapps.conf` file in `/etc/` and adjust the parameters by using the minimum acceptable snapshot size that allows to complete the backup without an error.

Before you begin

- Ensure that you have created the `/etc/guestapps.conf` file which doesn't exist by default.
- Ensure that the space requirements for Db2, MongoDB, and SAP HANA are met. For more information, see [“Space requirements for Db2 protection” on page 291](#), [“Space prerequisites for MongoDB protection” on page 353](#), and [“Space requirements for SAP HANA protection” on page 404](#).
- Ensure that you determined the minimum acceptable snapshot size. For instructions, see [“Determining the minimum acceptable snapshot size” on page 474](#).

About this task

The procedure for configuring the parameters in the `guestapps.conf` file is the same for Db2, MongoDB, and SAP HANA.

To understand how to allocate space in a volume group, consider the following example.

Assume that you have a volume group with total space capacity of 20 GB. The default value indicates that the minimum free space allocation is 2 GB (10%), and the maximum allocation is 5 GB (25%). During an inventory, the agent checks that the volume group has at least 2 GB of free space. Otherwise, an error message is displayed. If the free space is less than 5 GB, all available free space will be used to create the snapshot.



Figure 27. Space allocation in a volume group

Based on the previous example, you would configure the `guestapps.conf` file.

The following figures show the parameters in the `guestapps.conf` file for different databases. The shown values represent the defaults that are applied in the absence of the file.

Db2:

```
[DEFAULT]
```

```
Db2MinimumFreespaceInPercent = 10
```

```
Db2MaximumAllocationInPercent = 25
Db2MinimumSnapshotVolumeSize = 50
```

MongoDB:

```
[DEFAULT]
MongoMinimumFreespaceInPercent = 10
Db2MaximumAllocationInPercent = 25
Db2MinimumSnapshotVolumeSize = 50
```

SAP HANA:

```
[DEFAULT]
HANAMinimumFreespaceInPercent = 10
HANAMaximumAllocationInPercent = 25
HANAMinimumSnapshotVolumeSize = 50
```

Tip: When this logic is applied to thin-provisioned source volumes, the actual size is considered rather than the virtual size.

Procedure

To configure the `guestapps.conf` file, complete the following steps:

Remember: This procedure is an example for configuring the parameters for Db2.

1. The **Db2MinimumFreespaceInPercent** parameter, which specifies the minimum percentage of free space in a volume group that must be available to create snapshots. The agent will not start a backup operation if the percentage of free space in a volume group is less than this value.

To help ensure that the backup operation can take place successfully, you must add free space to the volume.

2. Calculate the value of the **Db2MaximumAllocationInPercent** parameter by using the following formula:

```
snapshot_size / sourceLV_size * 100
```

where *snapshot_size* is the size of the snapshot that you detected and *sourceLV_size* is the size of the source logical volume.

When backing up several logical volumes, use the maximum of all calculated **Db2MaximumAllocationInPercent** values.

3. Configure the parameters in `/etc/guestapps.conf` file and save the file.

What to do next

Repeat the backup operation. If issues recur, increase the **Db2MaximumAllocationInPercent** and optionally **Db2MinimumSnapshotVolumeSize** in the `guestapps.conf` file.

The **Db2MinimumSnapshotVolumeSize** parameter sets a lower limit for the snapshot size in absolute units (MB). In some cases, you might have to increase the **Db2MinimumSnapshotVolumeSize** in the `guestapps.conf` file for a successful backup. Then, repeat the backup operation.

Chapter 18. Product messages

IBM Spectrum Protect Plus components send messages with prefixes that help to identify which component they come from. Use the search option to find a particular message by using its unique identifier.

Messages consist of the following elements:

- A five-letter prefix.
- A number to identify the message.
- Message text that is displayed on screen and written to message logs.

Tip: Use your browser's search capability by using Ctrl+F to find the message code you are looking for.

The following example contains the Db2 agent prefix. When you click More, extra details that explain the reason for the message are shown.

```
Warning
Apr 16, 2019
9:14:37 AM
GTGGH0098
[myserver1.myplace.irl.ibm.com]
Database AC7 will not be backed up as it is ineligible for the backup operation. More
```

IBM Spectrum Protect Plus message prefixes

Messages have different prefixes to help you to identify the component that issues the message.

The following table identifies the prefix that is associated with each component.

Prefix	Component
CTGGA	IBM Spectrum Protect Plus
CTGGE	IBM Spectrum Protect Plus for Microsoft SQL Server
CTGGF	IBM Spectrum Protect Plus for Oracle
CTGGG	IBM Spectrum Protect Plus for Microsoft Exchange Server
CTGGH	IBM Spectrum Protect Plus for IBM Db2
CTGGI	IBM Spectrum Protect Plus for MongoDB
CTGGK	IBM Spectrum Protect Plus for Containers
CTGGL	IBM Spectrum Protect Plus for Amazon EC2
CTGGR	IBM Spectrum Protect Plus for Microsoft Office 365
CTGGS	IBM Spectrum Protect Plus for SAP HANA
CTGGT	IBM Spectrum Protect Plus for file systems

For a list of all messages, see IBM Documentation [here](#).

Appendix A. Search guidelines

Use filters to search for an entity such as a file or a restore point.

You can enter a character string to find objects with a name that exactly matches the character string. For example, searching for the term `string.txt` returns the exact match, `string.txt`.

Regular expression search entries are also supported. For more information, see [Search Text with Regular Expressions](#).

You can also include the following special characters in the search. You must use a backslash (`\`) escape character before any of the special characters:

```
+ - & | ! ( ) { } [ ] ^ " ~ * ? : \
```

For example, to search for the file `string[2].txt`, enter the `string\[2\].txt`.

Searching with wildcards

You can position wildcards at the beginning, middle, or end of a string, and combine them within a string.

Match a character string with an asterisk

The following examples show search text with an asterisk:

- `string*` searches for terms like `string`, `strings`, or `stringency`
- `str*ing` searches for terms like `string`, `straying`, or `straightening`
- `*string` searches for terms like `string` or `shoestring`

You can use multiple asterisk wildcards in a single text string, but multiple wildcards might considerably slow down a large search.

Match a single character with a question mark

The following examples show search text with a question mark:

- `string?` searches for terms like `strings`, `stringy`, or `string1`
- `st??ring` searches for terms like `starring` or `steering`
- `???string` searches for terms like `hamstring` or `bowstring`

Appendix B. Accessibility features for the IBM Spectrum Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

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

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

The product documentation in IBM Documentation is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

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

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

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

Vendor software

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

Related accessibility information

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

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

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

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

Trademarks

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

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 [IBM Privacy Policy](#) and IBM's Online Privacy Statement at [IBM Product Privacy](#) in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at [IBM Product Privacy](#).

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products. See the [IBM Spectrum Protect glossary](#).

Index

A

- Access control
 - MongoDB [353](#)
- accessibility features [483](#)
- ad hoc jobs
 - creating [439](#)
- Add Db2 partitions [292](#)
- adding
 - Amazon EC2 account [258](#)
 - Hyper-V servers [243](#)
 - identities [469](#)
 - LDAP server [162](#)
 - Oracle application servers [377](#)
 - sites [158](#)
 - SMTP server [164](#)
 - SQL Server application servers [390](#)
 - vCenter Server instances [213](#)
 - virtual disks to a vCenter virtual machine [182](#)
 - vSnap servers [40](#)
- Adding a filesystem [268](#)
- Adding Db2 [292](#)
- Adding MongoDB [355](#)
- Adding SAP HANA [406](#)
- Administrative Console, logging on to [180](#)
- Advanced backup options [48](#)
- Amazon EC2
 - accounts
 - adding [258](#)
 - backup job, creating [259](#)
 - detecting resources [259](#)
 - IAM user, creating [257](#)
- application server
 - Db2 [289](#), [404](#)
- AWS EC2
 - restore job, creating [261](#)

B

- Backing up
 - Db2 [296](#)
 - file system data [272](#)
- backup jobs
 - creating
 - Amazon EC2 [259](#)
 - Hyper-V [245](#)
 - IBM Spectrum Protect Plus [427](#)
 - Oracle [379](#)
 - SQL Server [392](#)
 - VMware [220](#)
 - excluding virtual disks from [250](#)
 - excluding VMDKs from [225](#)
 - rerunning
 - on demand [439](#)
 - starting
 - on demand [433](#)
 - on schedule [127](#), [192](#), [196](#), [198](#), [203](#), [205](#), [206](#)

- backup policies, *See* SLA policies
- backup storage
 - advanced options, managing [48](#)
 - storage options, managing disks [44](#)
 - storage options, managing partners [47](#)
- backup storage server
 - storage options, managing [45](#), [47](#)
- Beta program
 - advantages [xi](#)
 - overview [xi](#)

C

- certificate
 - adding [167](#)
 - deleting [167](#)
 - vCenter Server instances [215](#)
- cloud provider
 - deleting [144](#)
 - editing [144](#)
- cloud server
 - adding a Microsoft azure cloud resource [141](#)
 - adding an Amazon S3 [138](#)
 - adding an IBM Cloud Object Storage resource [140](#)
 - adding an s3 compatible cloud resource [143](#)
- cold-data-cache storage pool [147](#)
- Configuring advanced storage options [50](#)
- Configuring backup storage
 - storage options, adding disks [44](#)
- copying data to tape [147](#)
- creating
 - reports [451](#)
 - resource groups [457](#)
 - roles [463](#)
 - SLA policies [127](#), [192](#), [196](#), [198](#), [203](#), [205](#), [206](#)
 - users
 - individual [466](#)
 - LDAP group [466](#)
 - VADP proxies [123](#), [227](#)

D

- data copy to tape
 - configuring [147](#)
- data protection [153](#), [155](#)
- Db2
 - system requirements [22](#)
- Db2 log backup [299](#)
- DEFINE STGPOOL command [147](#)
- deleting
 - identities [470](#)
 - jobs [438](#)
 - LDAP server [165](#)
 - resource groups [461](#)
 - roles [465](#)
 - sites [162](#)
 - SLA policies [211](#)

- deleting (*continued*)
 - SMTP server [165](#)
 - users [468](#)
- Detailed process logs
 - Microsoft [365](#) [284](#)
- Detecting
 - Db2 [295](#)
 - file system resources [271](#)
- disability [483](#)

E

- early availability updates, obtaining and applying [93](#)
- editing
 - identities [470](#)
 - jobs and job schedules [437](#)
 - LDAP server [165](#)
 - resource groups [460](#)
 - roles [465](#)
 - settings [165](#)
 - sites [160](#)
 - SLA policies [210](#)
 - SMTP server [165](#)
 - users [467](#)
- efix [93](#)
- Exchange Server
 - system requirements [22](#)
- expire job session [429](#)

F

- fenced network, creating [239](#)
- file systems
 - system requirements [22](#)
- files
 - restoring [264](#)
 - searching for [481](#)
- Finding Db2 [295](#)
- Finding file system drives [271](#)
- firewalls [24](#)

G

- global preferences
 - configuring [173](#)

H

- Hyper-V
 - adding [243](#)
 - backup job, creating [245](#)
 - backup job, excluding virtual disks from an SLA policy [250](#)
 - installing on virtual appliance [32](#)
 - restore job, creating [251](#)
 - servers
 - detecting resources for [244](#)
 - enabling WinRM [244](#)
 - testing connection to [244](#)
 - virtual appliance
 - accessing [181](#)

I

- IBM Cloud Pak for Multicloud Management, integrating IBM Spectrum Protect Plus [19](#)
- IBM Knowledge Center [ix](#)
- IBM Spectrum Protect Operations Center
 - Accessing from IBM Spectrum Protect Plus [107](#)
 - adding IBM Spectrum Protect Plus to [131](#)
 - monitoring IBM Spectrum Protect Plus from [107](#), [135](#)
 - starting from IBM Spectrum Protect Plus [135](#)
 - URL, setting [134](#)
- IBM spectrum protect server
 - adding a repository server [156](#)
 - registering a repository server [156](#)
- identities
 - adding [469](#)
 - deleting [470](#)
 - editing [470](#)
- installing
 - as a virtual appliance [29](#)
 - download packages, obtaining [29](#)
 - post installation tasks [22](#)
 - virtual appliance
 - on Hyper-V [32](#)
 - on VMware [30](#)
 - vSnap servers
 - Hyper-V environment [37](#)
 - physical environment [35](#)
 - VMware environment [36](#)
- inventory
 - file systems [271](#)
- iSCSI utilities
 - installing [59](#)

J

- jobs
 - canceling [438](#)
 - concurrent, viewing [437](#)
 - creating [432](#)
 - deleting [438](#)
 - editing [437](#)
 - logs
 - downloading [437](#)
 - viewing [437](#)
 - names of [431](#)
 - pausing [437](#)
 - progress, viewing [436](#)
 - releasing [437](#)
 - rerunning [439](#)
 - schedules, editing [437](#)
 - starting
 - on demand [433](#)
 - on schedule [127](#), [192](#), [196](#), [198](#), [203](#), [205](#), [206](#)
 - types of [431](#)
 - viewing [434](#)
- Jobs and Operations [431](#)

K

- key
 - adding [166](#), [168](#)

key (*continued*)
deleting [167, 169](#)
keyboard [483](#)
keys and certificates
for IBM Spectrum Protect Plus user interface [171](#)
for secondary storage and resources [166, 170](#)
Knowledge Center [ix](#)

L

LDAP
group, creating a user account for [466](#)
server
adding [162](#)
deleting [165](#)
settings, editing [165](#)
Log archiving
Db2 [299](#)
SAP HANA [412](#)
logs
audit
downloading [452](#)
viewing [452](#)
system
downloading [471](#)
viewing [471](#)

M

message
prefixes [479](#)
messages [479](#)
Microsoft [365 281](#)
Microsoft 365 log files
Detailed [284](#)
MongoDB
system requirements [22](#)
MongoDB application server [352](#)

N

network
testing [171, 172](#)
Network configuration [45](#)
New in IBM Spectrum Protect Plus Version 10.1.13 [xiii](#)
NICs [45](#)

O

object client [153, 155](#)
Object Storage
Amazon S3 [138](#)
offline updates, virtual appliance [85](#)
online updates, virtual appliance [85](#)
Operations Center
Accessing from IBM Spectrum Protect Plus [107](#)
adding IBM Spectrum Protect Plus to [131](#)
monitoring IBM Spectrum Protect Plus from [107, 135](#)
starting from IBM Spectrum Protect Plus [135](#)
URL, setting [134](#)
Ops Manager
MongoDB [357](#)
Oracle

Oracle (*continued*)
application servers
adding [377](#)
detecting resources for [379](#)
testing connection to [379](#)
backup job, creating [379](#)
multithreaded databases [377](#)
restore job, creating [382](#)
system requirements [22](#)

P

password
superuser
changing [468](#)
preferences
global
configuring [173](#)
prerequisites
Db2 [289](#)
file systems [267](#)
MongoDB [352](#)
SAP HANA [404](#)
Prerequisites
MongoDB [353](#)
product overview [18](#)
publications [ix](#)

Q

quick start [95, 101, 104](#)

R

RBAC
MongoDB [353](#)
registering
vSnap servers [40](#)
repair vSnap [59](#)
Replication partners [47](#)
reports
custom, creating [451](#)
file systems [444](#)
Microsoft 365 [444](#)
running
on demand [450](#)
on schedule [452](#)
running VM [449](#)
types of
backup storage utilization [443](#)
protection [444](#)
system [447](#)
repository server provider
deleting [158](#)
editing [158](#)
rerunning
jobs
on demand [439](#)
Rescan
After expanding storage [45](#)
Rescan vSnap [45](#)
resource groups
creating [457](#)

resource groups (*continued*)

- deleting [461](#)
- editing [460](#)
- types of [458](#)

restore jobs

creating

- AWS EC2 [261](#)
- Hyper-V [251](#)
- IBM Spectrum Protect Plus [428](#)
- Oracle [382](#)
- SQL Server [396](#)
- VMware [232](#)

running

- AWS EC2 [261](#)
- Hyper-V [251](#)
- Oracle [382](#)
- SQL Server [396](#)
- VMware [232](#)

restore points, deleting [430](#)

restore points, managing [429](#)

Restoring

- Db2 [300](#), [305](#), [308](#)
- file system [277](#)
- SAP HANA [413](#), [418](#), [421](#)

Restoring Db2

- Alternate instance [308](#)
- Original instance [305](#)

Restoring SAP HANA

- Alternate instance [421](#)
- Original instance [418](#)

roles

- creating [463](#)
- deleting [465](#)
- editing [465](#)
- permission types [463](#)

S

SAP HANA

- system requirements [22](#)

SAP HANA log backup [412](#)

Schedule jobs

- Backup [297](#), [317](#), [360](#), [411](#)

scripts for backup and restore operations

- uploading [441](#)

service level agreement, *See* SLA policies

Setting Db2

- SLA options [298](#)

sites

- adding [158](#)
- deleting [162](#)
- editing [160](#)
- throttling [158](#), [160](#)

SLA [297](#), [317](#), [360](#), [411](#)

SLA options

- Db2 [298](#)

SLA policies

- adding [127](#), [192](#), [196](#), [198](#), [203](#), [205](#), [206](#)
- deleting [211](#)
- editing [210](#)

SMTP

server

- adding [164](#)
- deleting [165](#)

SMTP (*continued*)

server (*continued*)

- settings, editing [165](#)

snapshot retention [429](#)

sponsor user program

- advantages [xi](#)
- overview [xi](#)

SQL Server

application servers

- adding [390](#)
- detecting resources for [391](#)
- testing connection to [391](#)

backup job, creating [392](#)

requirements for data protection [389](#)

restore job, creating [396](#)

system requirements [22](#)

SSL certificate, uploading [171](#)

starting

IBM Spectrum Protect Plus [96](#)

jobs

- on demand [433](#)
- on schedule [127](#), [192](#), [196](#), [198](#), [203](#), [205](#), [206](#)

superuser

- changing password [468](#)
- changing password and name [468](#)

system requirements

- components [21](#)
- Db2 [22](#)
- Exchange Server [22](#)
- file index and restore [21](#)
- file systems [22](#)
- hypervisors [21](#)
- MongoDB [22](#)
- Oracle [22](#)
- SAP HANA [22](#)
- SQL Server [22](#)

T

testing

- vCenter Server user account [218](#)

Testing connection

- Db2 [296](#)

Testing connection file systems [271](#)

time zone, setting [182](#)

Transport Encryption [50](#)

U

Updating

- vSnap server [88](#)

user access [11](#), [455](#)

users

- deleting [468](#)
- editing [467](#)
- individual, creating [466](#)
- LDAP group, creating [466](#)
- resource groups
 - creating [457](#)
 - deleting [461](#)
 - editing [460](#)
 - types of [458](#)

roles

users (*continued*)
roles (*continued*)
 creating [463](#)
 deleting [465](#)
 editing [465](#)
 permission types [463](#)

V

VADP proxies
 creating [123](#), [227](#)
 options, setting [229](#)
 uninstalling [231](#)
 updating [91](#)

virtual appliance
 accessing
 in Hyper-V [181](#)
 in VMware [181](#)
 adding a disk to [183](#)
 adding storage capacity [183](#)
 installing
 on Hyper-V [32](#)
 on VMware [30](#)

Virtual appliance
 updating [85](#)

virtual environments [153](#), [155](#)

VMware
 backup job, creating [220](#)
 backup job, excluding VMDKs from SLA policy [225](#)
 installing on virtual appliance [30](#)
 restore job
 creating a fenced network [239](#)
 restore job, creating [232](#)
 vCenter Server instances
 adding [213](#)
 editing [215](#)
 vCenter Server user account
 testing [218](#)
 vCenter Server, detecting resources [219](#)
 vCenter Server, testing connection to [219](#)
 virtual appliance
 accessing [181](#)
 virtual machine privileges, required [216](#)

vSnap
 updating [89](#)

vSnap recovery [59](#)

vSnap server
 administering
 kernel headers
 kernel tools [67](#)
 network administration [65](#)
 storage administration [62](#)
 user administration [60](#)
 vSnap certificate [66](#)
 change throughput [58](#)
 editing [42](#)
 initializing
 advanced [54](#)
 simple [53](#)
 storage pools, expanding [58](#)
 Unregistering [43](#)

vSnap servers
 adding [40](#)
 installing

vSnap servers (*continued*)
 installing (*continued*)
 Hyper-V environment
 [37](#)
 physical environment [35](#)
 VMware environment [36](#)
 registering [40](#)
 uninstalling [39](#)

W

WinRM, enabling for connection to Hyper-V servers [244](#)



Product Number: 5737-F11