

IBM Spectrum Protect Plus
Version 10.1.5

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 395](#).

This edition applies to version 10, release 1, modification 5 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, 2019.**

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
What's new in Version 10.1.5.....	xi
Getting involved in product development.....	xiii
Sponsor user program.....	xiii
Beta program.....	xiii
Chapter 1. Product overview.....	1
Product components.....	1
Product dashboard.....	3
Alerts.....	4
Role-based access control.....	5
Replicate backup-storage data.....	5
Copy snapshots to secondary backup storage.....	6
IBM Spectrum Protect Plus on IBM Cloud.....	8
IBM Spectrum Protect Plus on AWS.....	9
Chapter 2. Installing IBM Spectrum Protect Plus.....	11
Product deployment roadmap.....	11
System requirements	11
Component requirements	11
Hypervisor requirements	26
File indexing and restore requirements.....	27
Microsoft Exchange Server requirements.....	31
Db2 requirements.....	36
MongoDB requirements.....	39
Office 365 requirements.....	41
Oracle requirements.....	44
Microsoft SQL Server requirements.....	49
Kubernetes Backup Support requirements.....	54
Obtaining the IBM Spectrum Protect Plus installation package.....	58
Installing IBM Spectrum Protect Plus as a VMware virtual appliance.....	58
Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance.....	60
Assigning a static IP address.....	62
Uploading the product key.....	62
Editing firewall ports.....	63
Installing iSCSI initiator utilities.....	65
Chapter 3. Installing vSnap servers.....	67
Installing a vSnap server.....	67
Installing a physical vSnap server.....	67
Installing a virtual vSnap server in a VMware environment.....	68
Installing a virtual vSnap server in a Hyper-V environment.....	69
Uninstalling a vSnap server.....	70
Chapter 4. Managing vSnap servers.....	73

Adding a vSnap server as a backup storage provider.....	73
Editing settings for a vSnap server.....	74
Configuring backup storage options.....	75
Initializing the vSnap server.....	80
Completing a simple initialization.....	80
Completing an advanced initialization.....	81
Expanding a vSnap storage pool.....	81
Establishing a replication partnership for vSnap servers.....	82
Changing the throughput rate.....	82
Replacing a failed vSnap server.....	83
vSnap server administration reference	84
Storage management.....	84
Network management.....	87
Kernel headers and tools.....	87
Chapter 5. Getting off to a quick start.....	89
Start IBM Spectrum Protect Plus.....	91
Manage sites.....	92
Create backup policies.....	93
Create a user account for the application administrator.....	95
Add resources to protect.....	96
Add resources to a job definition.....	98
Start a backup job.....	99
Run a report.....	100
Chapter 6. Updating IBM Spectrum Protect Plus components.....	103
Updating the IBM Spectrum Protect Plus virtual appliance.....	103
Additional steps for updating virtual machines in Hyper-V Replica environments.....	105
Updating vSnap servers.....	105
Updating the operating system for a physical vSnap server.....	106
Updating the operating system for a virtual vSnap server.....	106
Updating a vSnap server.....	107
Updating VADP proxies.....	108
Applying early availability updates.....	109
Chapter 7. Configuring the system environment.....	111
Managing secondary backup storage.....	111
Managing cloud storage.....	111
Managing repository server storage.....	117
Managing keys and certificates.....	121
Managing sites.....	124
Adding a site.....	124
Editing a site.....	125
Deleting a site.....	126
Managing LDAP and SMTP servers.....	126
Adding an LDAP server.....	127
Adding an SMTP server.....	128
Editing settings for an LDAP or SMTP server.....	129
Deleting an LDAP or SMTP server.....	129
Configuring global preferences.....	130
Customizing global preferences.....	133
Logging on to the administrative console.....	135
Setting the time zone.....	136
Uploading an SSL certificate from the administrative console.....	137
Uploading an SSL certificate from the command line.....	138
Logging on to the virtual appliance.....	138
Accessing the virtual appliance in VMware.....	138

Accessing the virtual appliance in Hyper-V.....	139
Testing network connectivity.....	139
Running the Service Tool from a command line.....	139
Running the Service Tool remotely.....	140
Adding virtual disks.....	141
Adding a disk to the virtual appliance.....	141
Adding storage capacity from a new disk to the appliance volume.....	142
Chapter 8. Managing SLA policies for backup operations.....	145
Creating an SLA policy.....	145
Editing an SLA policy.....	149
Deleting an SLA policy.....	149
Chapter 9. Protecting hypervisors.....	151
VMware.....	151
Adding a vCenter Server instance.....	151
Backing up VMware data.....	155
Managing VADP backup proxies.....	160
Restoring VMware data.....	164
Hyper-V.....	174
Adding a Hyper-V server.....	174
Backing up Hyper-V data.....	176
Restoring Hyper-V data.....	179
Restoring files.....	185
Chapter 10. Protecting applications.....	189
Db2.....	189
Prerequisites for Db2.....	189
Adding a Db2 application server.....	192
Backing up Db2 data.....	196
Restoring Db2 data	202
Exchange Server.....	214
Prerequisites.....	214
Privileges	214
Adding an Exchange application server.....	216
Backing up Exchange databases.....	217
Incremental forever backup strategy.....	220
Restoring Exchange databases.....	220
Accessing Exchange database files with instant access mode.....	248
MongoDB.....	252
Prerequisites for MongoDB.....	252
Adding a MongoDB application server.....	254
Backing up MongoDB data.....	259
Restoring MongoDB data	263
Microsoft Office 365.....	278
Registering with Azure Active Directory	278
Registering the Office 365 tenant with IBM Spectrum Protect Plus	279
Backing up Office 365 data.....	280
Restoring Office 365 data.....	281
Oracle.....	282
Adding an Oracle application server.....	282
Backing up Oracle data.....	284
Restoring Oracle data.....	287
SQL Server.....	293
Adding an SQL Server application server.....	294
Backing up SQL Server data.....	296
Restoring SQL Server data.....	300

Chapter 11. Protecting containers.....	307
Overview.....	307
Backup and restore types.....	308
SLA policies.....	309
User roles.....	309
Kubernetes Backup Support requests.....	310
Security features.....	311
Installing Kubernetes Backup Support.....	313
Prerequisites.....	313
Installing and deploying Kubernetes Backup Support.....	315
Uninstalling Kubernetes Backup Support.....	320
Completely uninstalling Kubernetes Backup Support.....	321
Backing up container data.....	322
Scheduling backups of persistent volumes.....	322
Backing up a persistent volume on demand.....	324
Restoring container data.....	326
Managing container backup and restore jobs.....	329
Viewing the status of backup and restore jobs.....	329
Pausing scheduled backups.....	333
Resuming scheduled backups.....	334
Deleting container backups.....	334
Viewing jobs and running reports.....	336
Viewing job logs.....	336
Creating backup history reports.....	337
 Chapter 12. Protecting IBM Spectrum Protect Plus.....	 339
Backing up the application.....	339
Restoring the application.....	339
Managing restore points.....	340
Expiring job sessions.....	340
Deleting resource metadata from the catalog.....	341
 Chapter 13. Managing jobs and operations.....	 343
Job types.....	343
Concurrent jobs.....	344
Creating jobs and job schedules.....	344
Starting jobs on demand.....	345
Viewing jobs and job logs.....	346
Pausing and resuming jobs.....	347
Editing jobs and job schedules.....	347
Canceling jobs.....	348
Deleting jobs.....	348
Rerunning partially completed backup jobs.....	349
Running an ad hoc backup job.....	349
Configuring scripts for backup and restore operations.....	350
Uploading a script.....	350
Adding a script to a server.....	351
 Chapter 14. Managing reports and logs.....	 353
Types of reports.....	353
Backup storage utilization reports.....	353
Protection reports.....	354
System reports.....	356
Running VM environment reports.....	357
Report actions.....	359
Running a report.....	359

Creating a custom report.....	360
Scheduling a report.....	360
Collecting and reviewing audit logs for actions.....	361
Enabling the display of alerts in the system log.....	361

Chapter 15. Managing user access..... 363

Managing user resource groups.....	364
Creating a resource group.....	364
Editing a resource group.....	366
Deleting a resource group.....	367
Managing roles.....	367
Creating a role.....	368
Editing a role.....	370
Deleting a role.....	371
Managing user accounts.....	371
Creating a user account for an individual user.....	371
Creating a user account for an LDAP group.....	371
Editing user account credentials.....	372
Deleting a user account.....	373
Managing identities.....	373
Adding an identity.....	373
Editing an identity.....	373
Deleting an identity.....	374

Chapter 16. Licensing overview.....375

Software License Metric (SLM) tags.....	375
Integration with IBM License Metric Tool (ILMT).....	376

Chapter 17. Troubleshooting.....377

Collecting log files for troubleshooting.....	377
Troubleshooting Kubernetes Backup Support.....	377
Quick reference.....	377
Troubleshooting backups and restores.....	379
Collecting Kubernetes Backup Support log files.....	385
Setting the trace level of log files.....	386
Viewing trace logs for Kubernetes Backup Support.....	387

Chapter 18. Product messages.....389

Message prefixes.....	389
-----------------------	-----

Appendix A. Search guidelines.....391

Appendix B. Accessibility.....393

Notices.....395

Glossary.....399

Index..... 401

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.

In this publication, it is assumed that you have an understanding of the applications that support IBM Spectrum Protect Plus as described in [“System requirements”](#) on page 11.

Publications

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

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

What's new in Version 10.1.5

IBM Spectrum Protect Plus Version 10.1.5 introduces new features and updates.

For a list of new features and updates in this release and 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.

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](#)

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

For more information about the beta program, contact Mary Anne Filosa at <mailto:mfilosa@us.ibm.com>.

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

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.

Product components

The IBM Spectrum Protect Plus solution is provided as a self-contained 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 service in a deployment, even if the deployment is spread across multiple locations.

The IBM Spectrum Protect Plus server contains an onboard vSnap server and VMware vStorage API for Data Protection (VADP) proxy server. For smaller backup environments, these servers might be sufficient. However, for larger environments, more servers might be required.

The onboard vSnap server can be used to back up and restore a small number of virtual machines and evaluate IBM Spectrum Protect Plus operations. As your requirements for backing up and restoring data grow, your vSnap storage can be expanded by adding external vSnap servers. By adding external vSnap servers to your environment, you can reduce the load on the IBM Spectrum Protect Plus appliance.

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 preferred method is to localize data movement to sites by placing vSnap servers and VADP proxies together at a single site. The placement of backup data to a site is governed by service level agreement (SLA) policies.

vSnap server

This component is a pool of disk storage that receives data from production systems for the purposes of data protection or reuse. The vSnap server consists of one or more disks and can be scaled up (adding disks to increase capacity) or scaled out (introducing multiple vSnap servers to increase overall performance). 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 3](#).

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

Item	Description
Alerts icon 	This icon opens the Alerts window. For more information about alerts, see “Alerts” on page 4 .
Help icon 	This icon opens the online help system.
User menu	This menu shows the name of the user who is logged on. The menu provides access to product information and documentation, logs, and the user sign out option.

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 admin privileges.

Administrative console

The administrative console is used to install software patches and updates and to complete other administrative tasks such as managing security certificates, starting and stopping IBM Spectrum Protect Plus, and changing the time zone for the application.

Example 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

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 hypervisor resources, 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 belong to a role that has the required permissions. Ensure that your user account belongs to a role that has the required permissions before you start the task.

To set up and manage user access, see [Chapter 15, “Managing user access,” on page 363](#).

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 145.

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:

- 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.

Copy snapshots to secondary backup storage

The vSnap server is the primary backup location for snapshots. All IBM Spectrum Protect Plus environments have at least one vSnap server. Optionally, you can copy snapshots from a vSnap server to secondary backup storage.

Terminology change: In previous releases, the process of copying data from IBM Spectrum Protect Plus to secondary backup storage was known as *offloading* data. Beginning with IBM Spectrum Protect Plus Version 10.1.5, the process is known as *copying* data.

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, a full copy is created during the first copy operation. Subsequent copies are incremental and capture cumulative changes since the last copy operation.

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.

For information about how snapshot data is copied to standard object storage and archive object storage for each cloud storage system, see [“Cloud requirements” on page 23](#).

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 overview” on page 117 and “Adding a repository server as a backup storage provider” on page 119 .
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 111• “Adding IBM Cloud Object Storage as a backup storage provider” on page 112• “Adding Microsoft Azure cloud storage as a backup storage provider” on page 114• “Adding a repository server as a backup storage provider” on page 119
Create a backup policy that includes the storage.	See “Create backup policies” on page 93 .

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.

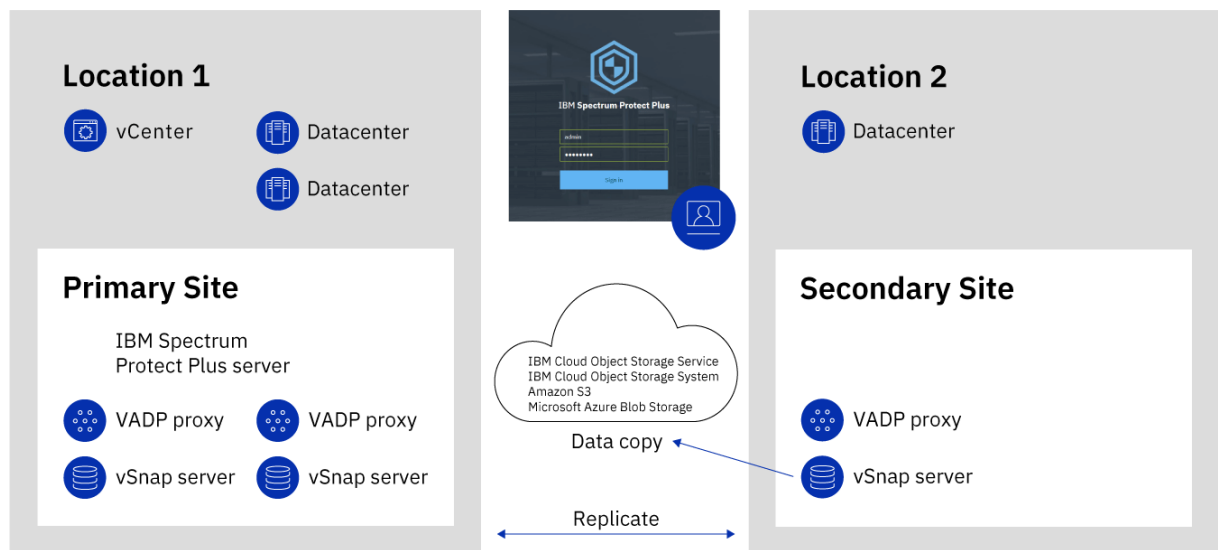


Figure 2. IBM Spectrum Protect Plus deployment across two geographical locations with copy to cloud storage

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.

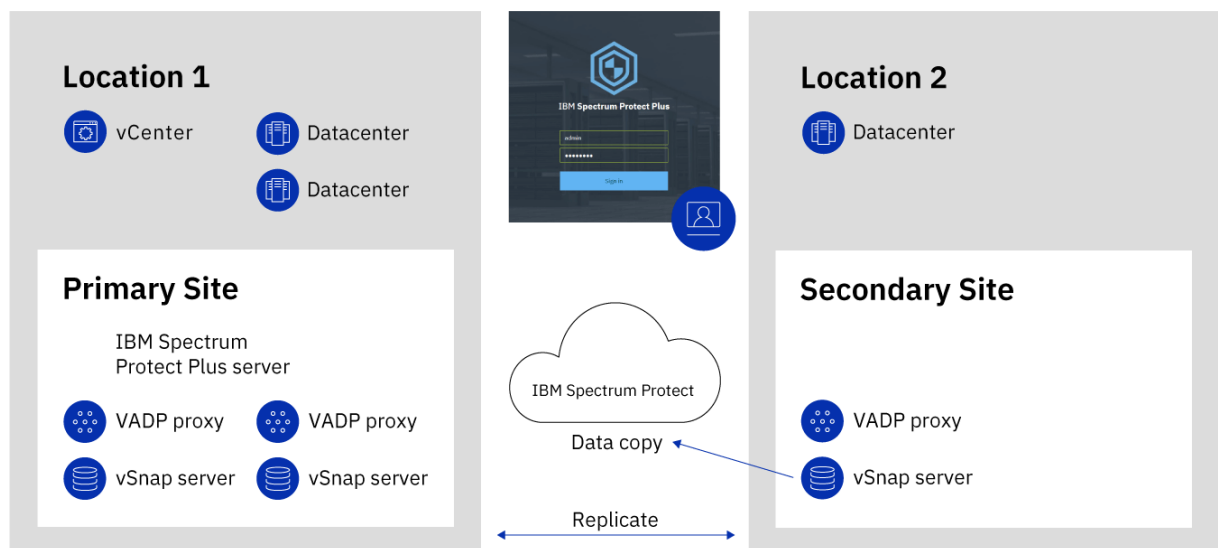


Figure 3. IBM Spectrum Protect Plus deployment across two geographical locations with copy to IBM Spectrum Protect

IBM Spectrum Protect Plus on IBM Cloud

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

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.

Available features in IBM Spectrum Protect Plus on IBM Cloud

IBM Spectrum Protect Plus supports both VMware and Microsoft Hyper-V environments.

However, IBM Spectrum Protect Plus on IBM Cloud supports only VMware 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. To find the documentation for the version of IBM Spectrum Protect Plus on IBM Cloud that you are using, go to the [online product documentation](#) and select the product version.

For more information

For information about how to order, install, and configure IBM Spectrum Protect Plus on IBM Cloud, see the following documentation. An IBMid is required to access the documentation.

- [Getting started with IBM Cloud for VMware Solutions](#)
- [Components and considerations for IBM Spectrum Protect Plus on IBM Cloud](#)
- [Managing IBM Spectrum Protect Plus on IBM Cloud](#)

IBM Spectrum Protect Plus on the AWS cloud platform

IBM Spectrum Protect Plus on the Amazon Web Services (AWS) cloud platform is a solution for users who run IBM Spectrum Protect Plus on premises, but want to protect databases that are running on the AWS cloud.

IBM Spectrum Protect Plus on AWS is a hybrid solution in which the IBM Spectrum Protect Plus server is on-premises and the vSnap server is on AWS.

The policy, system administration, access control, and other features of IBM Spectrum Protect Plus are managed and maintained by the on-premises IBM Spectrum Protect Plus server. Data from databases that are on AWS is then stored in the vSnap server that is also on AWS.

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 vSnap server to 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 10.1.3. If you want to use the current version of IBM Spectrum Protect Plus, follow the instructions in [Chapter 6, “Updating IBM Spectrum Protect Plus components,”](#) on [page 103](#) to complete an upgrade.

Chapter 2. Installing IBM Spectrum Protect Plus

Before you install IBM Spectrum Protect Plus, review the system requirements and installation procedures.

Product deployment roadmap

Follow the roadmap to install, configure, and start using IBM Spectrum Protect Plus.

Action	How to
Ensure that your system environment meets the hardware and software requirements.	See “System requirements ” on page 11.
Determine how to size, build, and place the components in your IBM Spectrum Protect Plus environment.	See the IBM Spectrum Protect Plus Blueprints.
Install IBM Spectrum Protect Plus.	See Chapter 2, “Installing IBM Spectrum Protect Plus,” on page 11.
If additional vSnap servers are required to support your environment, install and configure the servers.	See Chapter 3, “Installing vSnap servers,” on page 67.
If additional VMware vStorage API for Data Protection (VADP) proxies are required to support your environment, create and configure the proxies.	See “Managing VADP backup proxies” on page 160.
Complete the basic steps to set up and start using IBM Spectrum Protect Plus.	See Chapter 5, “Getting off to a quick start,” on page 89.

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.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

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

Ensure that you have the required system configuration and a supported browser to deploy and run IBM Spectrum Protect Plus.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

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.

Virtual machine installation

IBM Spectrum Protect Plus is installed as a virtual appliance. Before you deploy IBM Spectrum Protect Plus to the host, ensure that one of the following requirements is met:

- vSphere 6.0, 6.5, or 6.7
- Microsoft Hyper-V 2016 or Microsoft Hyper-V 2019

For initial deployment, configure your virtual appliance to meet the following minimum requirements:

- 64-bit 8-core machine
- 48 GB memory
- 536 GB disk storage for the virtual machine (VM)

Use a Network Time Protocol (NTP) server to synchronize the time zones across IBM Spectrum Protect Plus resources in your environment, such as the IBM Spectrum Protect Plus appliance, storage arrays, hypervisors, and application servers. If the clocks on the various systems are significantly out of sync, you might experience errors during application registration, metadata cataloging, inventory, backup, or file restore jobs. For more information about identifying and resolving timer drift, see the following VMware knowledge base article: [Time in virtual machine drifts due to hardware timer drift](#)

Browser support

Run IBM Spectrum Protect Plus from a computer that has access to the installed virtual appliance. IBM Spectrum Protect Plus was tested and certified with the following web browsers:

- Firefox 55.0.3 and later
- Google Chrome 60.0.3112 and later
- Microsoft Edge 40.15063 or Microsoft EdgeHTML 15.15063 and later

If your screen resolution is lower than 1024 x 768, some items might not fit in the window. Enable pop-up windows in your browser to access the help system and some IBM Spectrum Protect Plus operations.

IBM Spectrum Protect requirements

If you plan to use IBM Spectrum Protect as a repository server for copying data to cloud storage, ensure that you are using IBM Spectrum Protect V8.1.9.

IBM Spectrum Protect Plus ports

IBM Spectrum Protect Plus and associated services use the following ports. The ports use secure connections (HTTPS or SSL).

Note: In IBM Spectrum Protect Plus V10.1.3, port 9090 was used for online help. Starting with V10.1.4, this port is no longer required for online help. No further action is required.

Table 1. Communication ports when the target is an IBM Spectrum Protect Plus appliance				
Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	vSnap server	IBM Spectrum Protect Plus appliance	Provides access to troubleshoot and maintain tasks on the IBM Spectrum Protect Plus appliance by using the Secure Shell (SSH) protocol.

Table 1. Communication ports when the target is an IBM Spectrum Protect Plus appliance (continued)

Port	Protocol	Initiator	Target	Description
443	TCP	IBM Spectrum Protect Plus user interface	IBM Spectrum Protect Plus appliance	Provides web access by using the Hypertext Transfer Protocol Secure (HTTPS). This is the main entry point for client connections, which use the Secure Sockets Layer (SSL) protocol.
5671	TCP and Advanced Message Queuing Protocol (AMQP)	VMware vStorage API for Data Protection proxy (VADP proxy) host	IBM Spectrum Protect Plus appliance	Used to manage messages produced and consumed by the VADP proxy and VMware job management workers. This is a RabbitMQ message framework, which also facilitates job log management.
8090	TCP	Administrative console	IBM Spectrum Protect Plus appliance	Provides access for system administration. This extensible framework supports plugins that run operations such as system and network updates.
8761	TCP	VADP proxy host	IBM Spectrum Protect Plus appliance	Automatically discovers VADP proxies and is used by IBM Spectrum Protect Plus virtual machine (VM) backup operations.
111	TCP	vSnap server	IBM Spectrum Protect Plus appliance: Onboard vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communicating with ONC servers.

Table 1. Communication ports when the target is an IBM Spectrum Protect Plus appliance (continued)

Port	Protocol	Initiator	Target	Description
2049	TCP	vSnap server	IBM Spectrum Protect Plus appliance: Onboard vSnap server	Used to transfer Network File System (NFS) file sharing by the vSnap server.
3260	TCP	vSnap server	IBM Spectrum Protect Plus appliance: Onboard vSnap server	Used for Internet Small Computer System Interface (iSCSI) data transfer by the vSnap server.
20048	TCP	vSnap server	IBM Spectrum Protect Plus appliance: Onboard vSnap server	Used for NFS data transfer by the vSnap server.

Table 2. Communication ports when the initiator is an IBM Spectrum Protect Plus appliance

Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus appliance	vSnap server or VADP proxy host	Provides access to troubleshoot and maintain tasks on remote vSnap servers and the VADP proxy by using Secure Shell (SSH) protocol.
25	TCP	IBM Spectrum Protect Plus appliance	Email server that can be accessed by using the Simple Mail Transfer Protocol (SMTP)	Provides access to an email service.
389	TCP	IBM Spectrum Protect Plus appliance	Lightweight Directory Access Protocol (LDAP) server	Provides access to Active Directory Services.
443	TCP	IBM Spectrum Protect Plus appliance	Hypervisor: VMware ESXi host and vCenter	Provides access to ESXi and vCenter for managing operations.
636	TCP	IBM Spectrum Protect Plus appliance	LDAP server	Provides access to Active Directory Services by using the SSL protocol.

Table 2. Communication ports when the initiator is an IBM Spectrum Protect Plus appliance (continued)

Port	Protocol	Initiator	Target	Description
902	TCP	IBM Spectrum Protect Plus appliance	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components. By default, ESXi uses NFC for operations such as copying and moving data between datastores.
5985	TCP	IBM Spectrum Protect Plus appliance	Hypervisor: Hyper-V or agents that use the iSCSI initiator	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus appliance	Hypervisor: Hyper-V or agents that use the iSCSI initiator	Provides access to the Secure WinRM service for Windows-based servers.
8098	TCP	IBM Spectrum Protect Plus appliance	VADP proxy host	Supports Representational State Transfer application programming interface (REST API) communications between the IBM Spectrum Protect Plus appliance and the VADP proxy by using the Transport Layer Security (TLS) protocol.
8900	TCP	IBM Spectrum Protect Plus appliance	vSnap server	Supports the Open Virtual Appliance (OVA) or Installer version of the intelligent storage framework that is used as a target for data protection operations.

Use the following diagram as guidance for the communication paths managed by IBM Spectrum Protect Plus. This diagram can provide assistance for troubleshooting and network configuration for deployment scenarios.

- The labeled resources in the gray background represent the core services of the IBM Spectrum Protect Plus virtual appliance.
- The colors of the various modules represent different types of services as defined by the key, which is below the diagram.
- The area that is labeled **Firewall** represents the network firewall.
- Services that appear in the **Firewall** area indicative of the ports that are open on the firewall.
- Dashed arrows represent communication among resources and services.
- The arrow flows TOWARD the listening port.
- The port numbers that need to be open are indicated by the LISTENING port. For example, the vSnap service is represented as being external to the IBM Spectrum Protect Plus virtual appliance. It is listening on port 8900 as well as other ports.
- A component in the virtual appliance establishes a communication path with a connection to the vSnap service at port 8900.

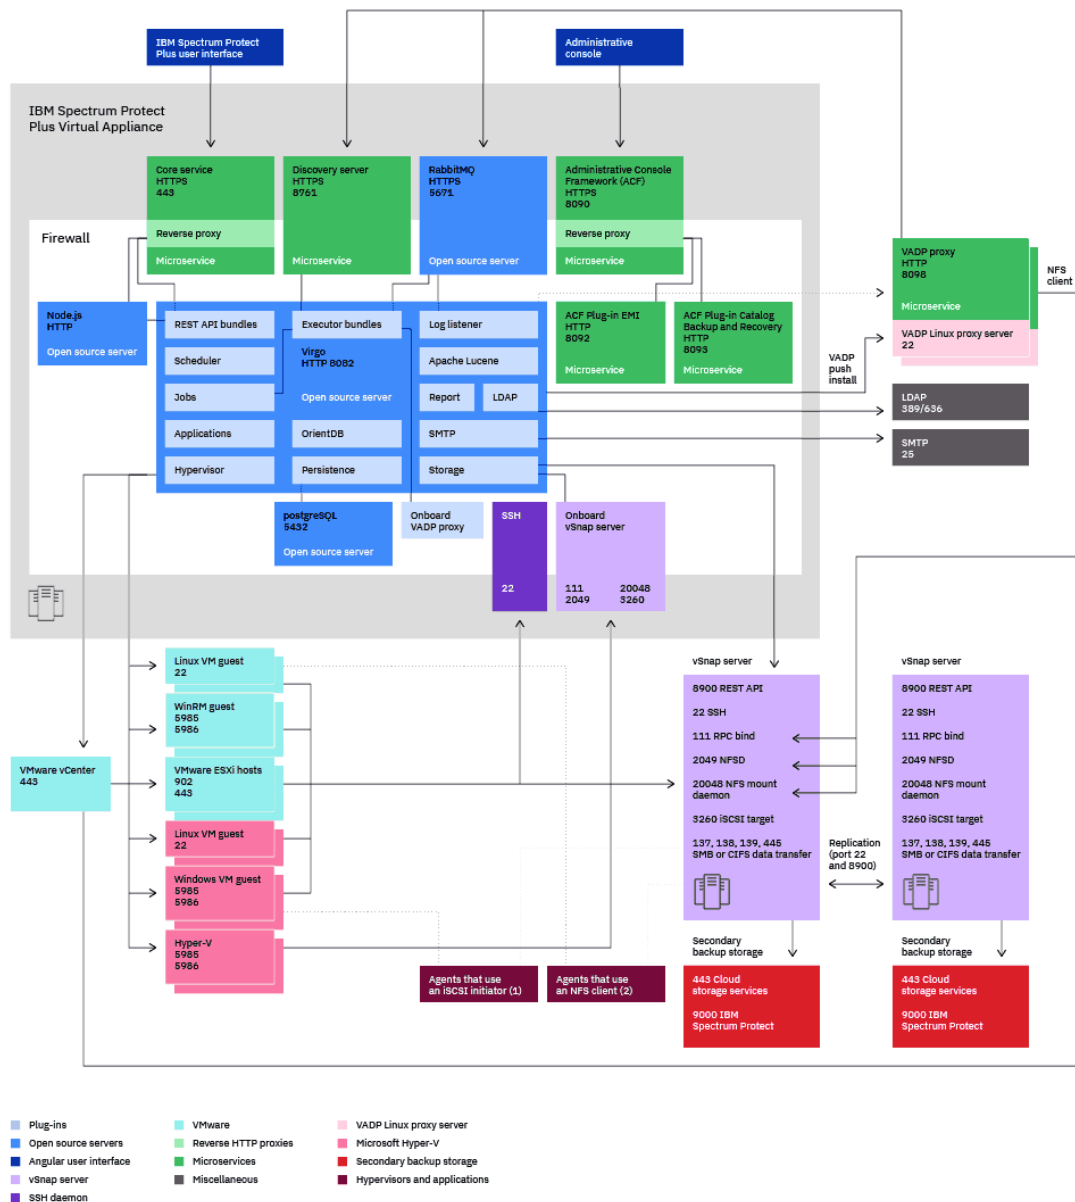


Figure 4. IBM Spectrum Protect Plus diagram

1. The following agents use an iSCSI initiator: Microsoft SQL Server and Microsoft Exchange.
2. The following agents use an NFS client: VMware, Microsoft Hyper-V, Db2®, MongoDB, Kubernetes, and Office 365.

Note: The IBM Spectrum Protect Plus virtual appliance contains the base components: IBM Spectrum Protect Plus server, site, vSnap server, vSnap pool, and VADP proxy. In the diagram, "Open source server" refers to the IBM Spectrum Protect Plus server.

vSnap server requirements

A vSnap server is the primary backup destination for IBM Spectrum Protect Plus. In either a VMware or Hyper-V environment, one vSnap server with the name localhost is automatically installed when the IBM Spectrum Protect Plus appliance is initially deployed. In larger backup enterprise environments, more vSnap servers might be required.

Allocate memory based on backup capacity for more efficient data deduplication. For more information about how to build an IBM Spectrum Protect Plus solution, see the [IBM Spectrum Protect Plus Blueprints](#).

For initial deployment, ensure that your VM or physical Linux server meet the following minimum requirements:

- 64-bit 8-core processor
- 32 GB memory
- 16 GB free space on root file system
- 128 GB free space in a separate file system mounted at `/opt/vsnap-data`

The Linux Network Management service must be installed and running.

Optionally, use a solid-state drive (SSD) to help improve backup and restore performance:

- To improve backup performance, configure the pool to use one or more log devices that are backed up to an SSD. Specify at least two log devices to create a mirrored log for better redundancy.
- To improve restore performance, configure the pool to use a cache device that is backed up to an SSD.

vSnap server VM installation requirements

Before you deploy the vSnap server to the host, ensure that one of the following requirements is met:

- vSphere 6.0, 6.5. or 6.7
- Microsoft Hyper-V 2016 or Microsoft Hyper-V 2019.

vSnap server physical installation requirements

Beginning with V10.1.3, IBM Spectrum Protect Plus provides new functions that requires the kernel levels that are supported in Red Hat Enterprise Linux (RHEL) 7.5 and CentOS 7.5. If you must use operating systems earlier than RHEL 7.5 and CentOS 7.5, use IBM Spectrum Protect Plus for physical vSnap V10.1.2 installations.

The following Linux operating systems are supported for IBM Spectrum Protect Plus V10.1.5 physical vSnap server installations:

- CentOS 7.1804 (7.5) (x86_64)
- CentOS 7.1810 (7.6) (x86_64)
- RHEL 7.5 (x86_64)
- RHEL 7.6 (x86_64)

If you are using the following operating systems, use IBM Spectrum Protect Plus for physical vSnap server V10.1.2 installations:

- CentOS 7.3.1611 (x86_64)
- CentOS 7.4.1708 (x86_64)
- RHEL 7.3 (x86_64)
- RHEL 7.4 (x86_64)

vSnap server ports

The following ports are used by vSnap servers. The ports use secure connections (HTTPS or SSL).

Table 3. Communication ports when the target is a vSnap server

Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus appliance, hypervisors or agents that use NFS client	vSnap server	Provides access to troubleshoot and maintain tasks on vSnap servers by using Secure Shell (SSH) protocol.
111	TCP	vSnap server, VADP proxy or agents that use the NFS client	vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communicating with ONC servers.
137	UDP	vSnap server or agents that use the ISCSI initiator	vSnap server	Provides a target port that is used by the vSnap Server Message Block (SMB) or the Common Internet File System (CIFS) to mount file system shares for transaction log backup and recovery operations.
138	UDP	vSnap server or agents that use the ISCSI initiator	vSnap server	Provides a target port that is used by the vSnap SMB or the CIFS to mount file system shares for transaction log backup and recovery operations.
139	TCP	vSnap server or agents that use the ISCSI initiator	vSnap server	Provides a target port that is used by the vSnap SMB or the CIFS to mount file system shares for transaction log backup and recovery operations.

Table 3. Communication ports when the target is a vSnap server (continued)

Port	Protocol	Initiator	Target	Description
445	TCP	vSnap server or agents that use the iSCSI initiator	vSnap server	Provides a target port that is used by the vSnap SMB or the CIFS to mount file system shares for transaction log backup and recovery operations.
2049	TCP	vSnap server, VADP proxy host or agents that use NFS client	vSnap server	Used to transfer Network File System (NFS) file sharing by the vSnap server.
3260	TCP	vSnap server or agents that use NFS client	vSnap server	Used for iSCSI data transfer by the vSnap servers.
8900	TCP	IBM Spectrum Protect Plus appliance	vSnap server	Supports the Open Virtual Appliance (OVA) or Installer version of the intelligent storage framework that is used as a target for data protection operations.
20048	TCP	vSnap server, VADP proxy host or agents that use NFS client	vSnap server	Mounts vSnap file systems on clients such as the VADP proxy, application servers, and virtualization data stores. This port is also used for NFS data transfer to vSnap servers.

VADP proxy requirements

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 your IBM Spectrum Protect Plus backup jobs. If proxies exist, the entire processing load is shifted from the IBM Spectrum Protect Plus appliance onto the proxies.

VADP proxies support the following VMware transport modes: File, SAN, HotAdd, NBDSSL, and NBD. For more information about VMware transport modes, see [Virtual Disk Transport Methods](#).

This feature is supported only in 64-bit quad core or higher configurations with a minimum kernel version of 2.6.32 in the following Linux environments:

- CentOS 6.5 and later maintenance and modification levels (beginning with 10.1.1 patch 1)
- CentOS 7.0 and later maintenance and modification levels (beginning with 10.1.1 patch 1)
- RHEL 6, Fix pack 4 and later maintenance and modification levels

- RHEL 7 and later maintenance and modification levels
- SUSE Linux Enterprise Server 12 and later maintenance and modification levels

For more information about how to build an IBM Spectrum Protect Plus solution, [IBM Spectrum Protect Plus Blueprints](#)

For initial deployment of a VADP proxy server, ensure that your Linux server meets the following minimum requirements:

- 64-bit quad core processor
- 8 GB random access memory (RAM) required, 16 GB preferred
- 60 GB free disk space

Because of increased CPU usage and concurrency on the VADP proxy server, the memory that is allocated on the proxy server must be increased. The proxy must be able to mount NFS file systems, which in many cases require an NFS client package to be installed. The package details vary based on the distribution.

Each proxy must have a fully qualified domain name and must be able to resolve and reach the vCenter. The vSnap servers must be reachable from the proxy.

Port 8098 on the VADP proxy server must be open when the proxy server firewall is enabled.

VADP proxy ports

The following ports are used by VADP proxies. The ports use secure connections (HTTPS or SSL).

<i>Table 4. Communication ports when the target is a VADP proxy host</i>				
Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus appliance	VADP proxy host	Provides access to troubleshoot and maintain tasks on VADP proxy hosts by using the SSH protocol.
8098	TCP	IBM Spectrum Protect Plus appliance	VADP proxy host	Supports Representational State Transfer application programming interface (REST API) communications between the IBM Spectrum Protect Plus appliance and the VADP proxy by using the Transport Layer Security (TLS) protocol.

Table 5. Communication ports when the initiator is a VADP proxy host

Port	Protocol	Initiator	Target	Description
111	TCP	VADP proxy host	vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communicating with ONC servers.
443	TCP	VADP proxy host	Hypervisor: VMware ESXi Host and vCenter	Provides access to ESXi and vCenter for managing operations.
902	TCP	VADP proxy host	Hypervisor: VMware ESXi Host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components. By default, ESXi uses NFC for operations such as copying and moving data between datastores.
2049	TCP	VADP proxy host	vSnap server	Used to transfer Network File System (NFS) file sharing by the vSnap server.
5671	TCP and AMQP	VADP proxy host	IBM Spectrum Protect Plus appliance	Used to manage messages produced and consumed by the VADP proxy and VMware job management workers. This is a RabbitMQ message framework, which also facilitates job log management.

Table 5. Communication ports when the initiator is a VADP proxy host (continued)				
Port	Protocol	Initiator	Target	Description
8761	TCP	VADP proxy host	IBM Spectrum Protect Plus appliance	Automatically discovers VADP proxies and is used by IBM Spectrum Protect Plus virtual machine (VM) backup operations.
20048	TCP	VADP proxy host	vSnap server	Mounts vSnap file systems on clients such as the VADP proxy, application servers, and virtualization data stores. This port is also used for NFS data transfer to vSnap servers.

Tip: VADP proxies can be pushed and installed on Linux-based servers over SSH port 22.

If the firewall command script is not available on your system, edit the firewall manually to add necessary ports, and restart the firewall. For instructions about editing firewall ports, see [“Editing firewall ports”](#) on page 63.

VADP proxy on vSnap server requirements

VADP proxies can be installed on the vSnap servers in your IBM Spectrum Protect Plus environment. A combination VADP proxy and vSnap server must meet the minimum requirements of both devices. Consider the system requirements of both devices and add the core and RAM requirements together to identify the minimum requirements of the combination VADP proxy and vSnap server. Ensure that your combination VADP proxy and vSnap server meet the following minimum requirements, which are the sum of the requirements for each device.

Ensure that your combination VADP proxy and vSnap server meet the following minimum requirements, which are the sum of the requirements for each device.

VADP proxy installed on a virtual vSnap server:

- 64-bit 8-core processor
- 48 GB RAM

All required VADP proxy and vSnap server ports must be open on the combination VADP proxy and vSnap server. Review the VADP proxy and vSnap ports sections of the system requirements for more information.

Cloud requirements

To copy data to cloud storage, ensure that your IBM Spectrum Protect Plus and cloud environments meet the following requirements.

Disk cache area

For all functions related to data copy and restore operations to and from cloud and archival targets, the vSnap server requires a disk cache area to be present on the vSnap server.

- During copy operations, this cache is used as a temporary staging area for objects that are pending upload to the cloud endpoint.

- During restore operations, the disk cache area is used to cache downloaded objects and to store any temporary data that might be written into the restore volume.

For instructions about sizing and installing the cache, see the [IBM Spectrum Protect Plus Blueprints](#).

Certificate requirements

- **Self-signed certificates:** If the cloud endpoint or repository server uses a self-signed certificate, you must specify certificate in Privacy Enhanced Mail (PEM) format when you register the cloud or repository server in the IBM Spectrum Protect Plus user interface.
- **Certificates signed by private Certificate Authority:** If the cloud endpoint or repository server uses a certificate signed by a private certificate authority (CA), the endpoint certificate must be specified (in PEM format) when you register the cloud or repository server in the IBM Spectrum Protect Plus user interface. In addition, you must add the root or intermediate certificate of the private CA to the system certificate store in each vSnap server by using the following procedure:
 1. Log in to the vSnap server console as the `serveradmin` user and upload any private CA certificates (in PEM format) to a temporary location.
 2. Copy each certificate file to the system certificate store directory (`/etc/pki/ca-trust/source/anchors/`) by running the following command:


```
$ sudo cp /tmp/private-ca-cert.pem /etc/pki/ca-trust/source/anchors/
```
 3. To incorporate the newly added custom certificate and update the system certificate bundle, run the following command:


```
$ sudo update-ca-trust
```
- **Certificates signed by public Certificate Authority:** If the cloud endpoint uses a public CA-signed certificate, no special action is required. The vSnap server validates the certificate by using the default system certificate store.

Network requirements

The following ports are used for communication between the vSnap servers and cloud or repository server endpoints.

Port	Protocol	Initiator	Target	Description
443	TCP	vSnap server	Cloud server endpoints	Allows the vSnap to communicate with Amazon Simple Storage Service (S3), Microsoft Azure, or IBM Cloud Object Storage endpoints.
9000	TCP	vSnap server	Repository server endpoints	Allows the vSnap to communicate with IBM Spectrum Protect (repository server) endpoints.

Any firewalls or network proxies that perform SSL Interception or Deep Packet Inspection for traffic between the vSnap servers and cloud endpoints might interfere with SSL certificate validation on vSnap servers. This interference can also cause cloud copy job failures. To prevent this interference,

the vSnap servers must be exempted from SSL interception and inspection in the firewall or proxy configuration.

Cloud provider requirements for standard and archive object storage

Native life-cycle management is not supported. IBM Spectrum Protect Plus manages the life-cycle of uploaded objects automatically by using an incremental-forever approach where older objects can still be used by newer snapshots. Automatic or manual expiration of objects outside of IBM Spectrum Protect Plus leads to data corruption.

If the cloud provider uses an SSL certificate that is self-signed or signed by a private certificate authority, see [Certificate requirements](#).

Amazon S3 cloud requirements

- **Standard Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket in one of the supported storage tiers must be specified: S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access, or S3 One Zone-Infrequent Access.
- **Archive Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket in one of the supported storage tiers must be specified: S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access, or S3 One Zone-Infrequent Access. IBM Spectrum Protect Plus directly uploads data files to the Glacier tier. Some small metadata files are stored in the default tier for the bucket. A copy of these metadata files is also placed into the Glacier tier for disaster recovery purposes.

IBM Cloud Object Storage requirements

- **Standard Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket must be specified. If the specified bucket has a WORM policy that locks objects for a certain time period, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock. The bucket must have the Name Index setting enabled.
- **Archive Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket must be specified. If the specified bucket has a WORM policy that locks objects for a certain time period, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock. IBM Spectrum Protect Plus creates a single life-cycle management rule on the bucket to migrate data files to the archive tier. The bucket must have the Name Index setting enabled.

Microsoft Azure requirements

- **Standard Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing container in a hot or cool storage account must be specified.
- **Archive Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing container in a hot or cool storage account must be specified. IBM Spectrum Protect Plus moves files between tiers on demand. Data files are immediately moved to the archive tier and temporarily returned to the hot tier only during restore operations. Some small metadata files are stored in the default tier for the container. A copy of these metadata files is also placed in the archive tier for disaster recovery purposes.

IBM Spectrum Protect (repository server) requirements

- **Standard Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, you cannot use an existing bucket. IBM Spectrum Protect Plus creates a uniquely named bucket for its own use.
- **Archive Object Storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, you cannot use an existing bucket. IBM Spectrum Protect Plus creates a uniquely named bucket for its own use. IBM Spectrum Protect Plus directly uploads data files to IBM Spectrum Protect tape storage. Some small metadata files are stored in IBM Spectrum Protect object storage. A copy of these metadata files is also placed on IBM Spectrum Protect tape storage for disaster recovery purposes.

Table 7. Copy and archive copy requirements for cloud providers		
Operation	Provider	Requirements
Copy	Amazon S3	An existing bucket must be specified from one of the supported storage tiers.
Copy	IBM Cloud Object Storage	An existing bucket must be specified. The bucket must have the Name Index setting enabled.
Copy	Microsoft Azure	An existing container must be specified from a hot or cool storage tier.
Copy	IBM Spectrum Protect	IBM Spectrum Protect Plus creates its own unique bucket.
Archive copy	Amazon S3	vSnap server must be able to communicate with IBM Spectrum Protect (repository server) endpoints.
Archive copy	IBM Cloud Object Storage	An existing bucket must be specified from the archive tier. The bucket must have the Name Index setting enabled.
Archive copy	Microsoft Azure	An existing container must be specified from the hot storage tier and archive tier.
Archive copy	IBM Spectrum Protect	IBM Spectrum Protect Plus creates its own unique bucket to be copied to IBM Spectrum Protect tape.

For more information about how to set up and copy data to specific cloud providers, see [Data offload to cloud object storage with IBM Spectrum Protect Plus](#).

Hypervisor requirements

Review the hypervisor requirements for IBM Spectrum Protect Plus.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

Hyper-V requirements

The Microsoft Hyper-V server must meet the following minimum requirements:

- Hyper-V Server 2016 or Microsoft Hyper-V on Windows Server 2016
- Hyper-V Server 2019 or Microsoft Hyper-V on Windows Server 2019

Backup and restore operations are only supported on virtual hard disks (VHDX). For more information, see [Known Issues and Limitations: IBM Spectrum Protect Plus V10.1.5.x](#)

IBM Spectrum Protect Plus protects virtual machines that are enabled to use the Hyper-V Replica feature. Depending on your Hyper-V environment, you might be required to update some SLA policies when you upgrade to IBM Spectrum Protect Plus V10.1.5. For more information about the upgrade requirements for

virtual machines in Hyper-V environments, see [“Additional steps for updating virtual machines in Hyper-V Replica environments”](#) on page 105.

The Microsoft iSCSI Initiator Service must be running on all Hyper-V servers, including cluster nodes. In the **Services** window, set the startup type for the Microsoft iSCSI Initiator Service to **Automatic** so that the service is available when the Hyper-V server or cluster node starts.

The **DiskPart** automount parameter must be enabled on the Hyper-V server. For more information about enabling the automount parameter, see the [Automount](#) topic on the Microsoft website.

Hyper-V servers can be registered by using a Domain Name System (DNS) name or an Internet Protocol (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 by DNS. If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus appliance by using the command line. If more than one Hyper-V server is set up in a cluster environment, you must add all of the servers to the `/etc/hosts` file. When you are registering the cluster in IBM Spectrum Protect Plus, register the Failover Cluster Manager.

VMware requirements

The following VMware vSphere versions are supported:

- vSphere 6.0, including all updates and patch levels
- vSphere 6.5, including all updates and patch levels
- vSphere 6.7, including all updates and patch levels

Ensure that the latest version of VMware Tools is installed in your environment.

Physical RDM (pRDM) volumes do not support snapshots. Virtual machines that contain one or more raw device-mapping (RDM) volumes that are provisioned in physical-compatibility mode (pRDM) are backed up. However, the pRDM volumes are not processed as part of the virtual machine backup operation.

IBM Spectrum Protect Plus V10.1.5 protect virtual machines managed by a VMware Cloud (VMC) on AWS Software-Defined Data Center (SDDC). For more information about this new capability, see [IBM Spectrum Protect Plus for VMware Cloud on AWS](#).

File indexing and restore requirements

Review file indexing and restore requirements for IBM Spectrum Protect Plus.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

iSCSI disks that are directly mapped to the guest operating system will not be indexed. Supported volumes include virtual machine disk (VMDK) or virtual hard disk (VHD) volumes that are mounted through the configuration of the associated virtual machine.

The amount of free space required for the metadata in the catalog depends on the total number of files present in the environment. To catalog 1 million files, the catalog volume in the IBM Spectrum Protect Plus appliance requires roughly 350 MB of free space per retained version. The space used by file indexing metadata is reclaimed when the corresponding backup instances expire.

VMware requirements

In the virtual machine settings under Advanced Configuration, the `disk.enableUUID` parameter must be set to `true`.

Windows requirements

Item	Description
Supported operating systems	<ul style="list-style-type: none">• Windows Server 2008 R2• Windows Server 2012 R2 and Windows Server 2012 R2 core• Windows Server 2016 and Windows Server 2016 Core• Windows Server 2019 and Windows Server 2019 core
Supported file systems	<ul style="list-style-type: none">• New Technology file system (NTFS)• Resilient file system (ReFS)• File allocation table (FAT)
Supported disk storage types	<p>Basic disks with the following partitions:</p> <ul style="list-style-type: none">• MBR (Master boot record)• GPT (GUID partition table) <p>Restriction: You cannot back up or restore files on dynamic disks.</p>

- IBM Spectrum Protect Plus supports only the operating systems that are available to your hypervisors. Review your hypervisor documentation for information about supported operating systems.
- File indexing and restore operations support SCSI disks in a Hyper-V environment. Integrated Drive Electronics (IDE) disks are not supported. Generation 1 virtual machines require IDE boot disks; however, if more SCSI disks are available, file indexing and restore operations are supported on those disks.
- Windows Remote Shell (WinRM) must be enabled.

Important: IBM Spectrum Protect Plus can protect and restore virtual machines with other file systems, but only the previously listed file systems are eligible for file indexing and restore.

- When files are indexed in a Windows environment, the following directories on the resource are skipped:

- \Program Files
 - \Program Files (x86)
 - \Windows
 - \winnt

Files within these directories are not added to the IBM Spectrum Protect Plus inventory and are not available for file recovery.

- Ensure that the latest version of VMware Tools is installed on VMware virtual machines, and Hyper-V Integration Services is installed on your Hyper-V virtual machines.
- File indexing and file restore of a Windows VM require that the Windows Powershell binary path is set in the %PATH% environment variable.
- 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 offloaded 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.

Space requirements

- The C:\ drive must have sufficient temporary space to save the file indexing results.
- When file systems are indexed, temporary metadata files are generated under the /tmp directory and are deleted when the indexing is complete. The amount of free space required for the metadata depends on the total number of files in the system. Ensure that there is approximately 350 MB of free space per 1 million files.

Connectivity requirements

- The host name of the IBM Spectrum Protect Plus appliance should be resolvable from the Windows virtual machine.
- The Internet Protocol (IP) address of the virtual machine selected for indexing must be visible to the vSphere client or Hyper-V Manager.
- The Windows virtual machine selected for indexing must allow outgoing connections to port 22 (SSH) on the IBM Spectrum Protect Plus appliance.
- All firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the server through WinRM.

Authentication and privilege requirements

The credentials that are specified for the virtual machine must include a user with the following privileges:

- The user identity must have the **Log on as a service** right, which is assigned through the Administrative Tools control panel on the local server (**Local Security Policy > Local Policies > User Rights Assignment > Log on as a service**).

For more information about the **Log on as a service** 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 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. Credentials must be established for the associated virtual machine through the **Guest OS Username** and **Guest OS Password** option within the associated backup job definition.
- The system login credential must have the permissions of the local administrator.

Kerberos requirements

- Kerberos-based authentication can be enabled through a configuration file on the IBM Spectrum Protect Plus appliance. This setting overrides the default Windows NTLM protocol. Kerberos does not allow local user accounts to be used and is suitable only for environments in which all virtual machines are on a single domain.
- For Kerberos-based authentication only, the user identity must be specified in the username@FQDN format. The specified user must be able to authenticate by using the registered password to obtain

a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain specified by the fully qualified domain name.

- Kerberos authentication also requires that the clock skew between the Domain Controller and the IBM Spectrum Protect Plus appliance is less than 5 minutes. The default Windows NTLM protocol is not time dependent.

Linux requirements

Item	Description
Supported operating systems	<ul style="list-style-type: none">• Red Hat Enterprise Linux (RHEL) 6.4 and later maintenance and modification levels• CentOS 6.4 and later maintenance and modification levels• RHEL 7.0 and later maintenance and modification levels• CentOS 7.0 and later maintenance and modification levels• SUSE Linux Enterprise Server 12.0 and later maintenance and modification levels
Supported file systems	<ul style="list-style-type: none">• ext2• ext3• ext4• XFS

- A file system created on a newer kernel version might not be mountable on a system with an older kernel, in which case restoring files from the newer to the older system is not supported.

IBM Spectrum Protect Plus supports only the operating systems available to your hypervisors. Review your hypervisor documentation for information about supported operating systems.

IBM Spectrum Protect Plus can protect and restore virtual machines with other file systems, but only the previously listed file systems are eligible for file indexing and restore operations.

- When file indexing is performed in a Linux environment, the following directories on the resource are skipped:

- /tmp
 - /usr/bin
 - /Drivers
 - /bin
 - /sbin

- Files in virtual file systems like /proc, /sys, and /dev are also skipped. Files within these directories are not added to the IBM Spectrum Protect Plus inventory and are not available for file recovery.

Space requirements

- The system disk must have sufficient temporary space to save the file indexing results.
- When file systems are indexed, temporary metadata files are generated under the /tmp directory and then deleted when the indexing is complete. The amount of free space required for the metadata depends on the total number of files in the system. Ensure that there is approximately 350 MB of free space per 1 million files.

Software requirements

- Red Hat Enterprise Linux / CentOS 6.x only: Ensure that the `util-linux-ng` package is current by running the following command: **yum update util-linux-ng**.

- Depending on your version or distribution, the package might be named `util-linux`.
- If data resides on LVM volumes, ensure that the LVM version is 2.0.2.118 or later. Run the **lvm version** command to check the version and run the **yum update lvm2** to update the package if necessary.
- If data resides on LVM volumes, the **lvm2-lvmetad** service must be disabled, as it can interfere with the ability of IBM Spectrum Protect Plus to mount and resignature volume group snapshots and clones. To disable the service, complete the following steps:

1. Run the following commands:

```
systemctl stop lvm2-lvmetad
systemctl disable lvm2-lvmetad
```

2. Edit `/etc/lvm/lvm.conf` and specify the following setting:

```
use_lvmetad = 0
```

For details of the **lvmetad** service, see [The Metadata Daemon \(lvmetad\)](#).

- If data resides on XFS file systems and the version of **xfspgrog**s is between 3.2.0 and 4.1.9, the file restore operation can fail due to a known issue in **xfspgrog**s that causes corruption of a clone or snapshot file system when its UUID is modified. To resolve this issue, update **xfspgrog**s to version 4.2.0 or later.

For more information, see [Debian Bug report logs](#).

Connectivity requirements

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 through Secure Shell (SSH). The secure file transfer protocol (SFTP) subsystem for SSH must also be enabled.

Authentication and privilege requirements

The credentials specified for the virtual machine must specify a user that has the following **sudo** privileges:

- The `sudoers` configuration must allow the user to run commands without a password.
- The `!requiretty` setting must be specified.

The recommended approach is to create a dedicated IBM Spectrum Protect Plus agent user with the following privileges. Sample configuration:

- Create user: `useradd -m sppagent`

where **sppagent** specifies the IBM Spectrum Protect Plus agent user.

- Set a password by using the command: `passwd <sppagent>`

Place the following lines at the end of your `sudoers` configuration file, typically `/etc/sudoers`. If your existing `sudoers` file is configured to import configurations from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory:

```
Defaults: sppagent !requiretty
sppagent ALL=(root) NOPASSWD:ALL
```

Microsoft Exchange Server requirements

Before you install IBM Spectrum Protect Plus, review the hardware and software requirements for the product and other components.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

The Exchange database backup and restore requirements for IBM Spectrum Protect Plus are as follows.

Configuration

Make sure that the Microsoft Exchange Server version that you are using is supported on your operating system.

Application Versions

- Microsoft Exchange Server 2013 CU16 and later CU and maintenance levels: Standard or Enterprise editions.
- Microsoft Exchange Server 2016 CU5 and later CU and maintenance levels: Standard and Enterprise editions.
- Microsoft Exchange Server 2019 and later maintenance levels: Standard and Enterprise editions.

Note: Microsoft Exchange database availability groups (DAG) are supported.

Operating Systems

- Windows Server 2012R2 and later maintenance levels (64-bit kernel): Standard and Datacenter editions
- Windows Server 2016 and later maintenance levels (64-bit kernel): Standard and Datacenter editions
- Windows Server 2019 and later maintenance levels (64-bit kernel): Standard and Datacenter editions

Note: Windows Server 2019 with the Server Core option is supported. However, the granular restore feature is not supported by the Server Core installation option.

Additional Notes

Install the latest Microsoft Exchange Server patches and updates in your environment.

IBM Spectrum Protect Plus supports Microsoft Exchange Server running on a physical (bare metal) server, also in a virtualization environment. The following virtualization environments are supported:

- VMware ESX guest operating system
- Microsoft Windows Hyper-V guest operating system

Incremental backups

IBM Spectrum Protect Plus uses update sequence number (USN) change journal technology to perform incremental backups in a Microsoft Exchange Server environment. The USN change journal provides write range tracking for a volume when the file size meets the minimum file size threshold requirement. The changed bytes offset and length extent information can be queried against a specific file.

To enable write range tracking, the system environment must meet the following requirements:

- Windows Server 2012 R2 or later
- New Technology File System (NTFS) Version 3.0 or later

The following technologies are not supported for changed bytes tracking:

- Resilient File System (ReFS)
- Server Message Block (SMB) 3.0 protocol
- SMB Transparent Failover (TFO)
- SMB 3.0 with scale-out file shares

By default, 512 MB of space is allocated for USN change journaling. In addition, when journal overflow is detected, the allocated space doubles in size when an overflow is detected, to a maximum of 2 GB.

The minimum space required for shadow copy storage is 100 MB, although more space might be required on systems with increased activity.

A base backup of a file is forced when the following conditions are detected:

- Journal discontinuity is reported. This can occur when the log reaches its maximum size, when journaling is disabled, or when the cataloged USN ID is changed.

- The file size is less than or equal to the tracking threshold size, which by default is 1 MB.
- A file is added after a previous backup job.

Software

Ensure that a supported version of a Windows 64-bit operating system is installed.

The following prerequisites from Microsoft are required and must be installed before you use IBM Spectrum Protect Plus:

- Windows PowerShell 4 or later
- Windows Management Framework 4 or later

When you use Microsoft Exchange Server 2013 and the granular restore feature, the minimum level that is supported for Microsoft Exchange Messaging API (MAPI) Client and Collaboration Data Objects (MAPI/CDO) is version 6.5.8320.0.

Note: MAPI and CDO are required for Microsoft Exchange Server 2013 only.

When you use the granular restore feature with Microsoft Exchange Server 2016 or 2019, Microsoft 32-bit Outlook 2013, Outlook 2016, or Outlook 2019 is required.

The following prerequisites from Microsoft are required, and are installed automatically by the IBM Spectrum Protect Plus granular restore feature, if not already present on your virtual machine.

- 32-bit Microsoft Visual C++ 2012 Redistributable Package
- 64-bit Microsoft Visual C++ 2012 Redistributable Package
- 32-bit Microsoft Visual C++ 2017 Redistributable Package
- 64-bit Microsoft Visual C++ 2017 Redistributable Package
- Microsoft .NET Framework 4.5
- Microsoft ReportViewer 2012 SP1 Redistributable Package
- Microsoft SQL Server 2012 System CLR Types
- Microsoft SQL Server 2014 System CLR Types
- Microsoft SQL Server 2016 System CLR Types

Tip: Installation of these prerequisites might require a system restart. To avoid a system restart, ensure that these prerequisites are installed before you start the IBM Spectrum Protect Plus granular restore feature.

Registration

Register each Microsoft Exchange Server with IBM Spectrum Protect Plus by name or IP address.

Restriction: The IP address must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. The fully qualified domain name of each Microsoft Exchange Server must be resolvable and can be routed from the IBM Spectrum Protect Plus server and from the vSnap server. The fully qualified domain name of the IBM Spectrum Protect Plus server must be resolvable and can be routed from the Microsoft Exchange servers.

The user identity must have sufficient rights to install and start the IBM Spectrum Protect Plus Tools Service on the node. These rights include **Log on as a service** rights. For more information, see the Microsoft article: [Add the Log on as a service Right to an Account](#)

Privileges

To use an Exchange database, an IBM Spectrum Protect Plus agent user must have the following permissions:

- Microsoft Exchange Server is protected by role-based authentication. For the Microsoft Exchange agent to work in your IBM Spectrum Protect Plus environment, you must set up the appropriate privileges.

- The Encrypting File System (EFS) must be enabled in the local or group domain policy, and a valid Domain Data Recovery Agent (DRA) certificate must be available.
- 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 Microsoft Exchange Server 2016 and 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.

For more information, see [“Privileges ”](#) on page 214.

Ports

The following ports are used by IBM Spectrum Protect Plus agents users. The ports use secure connections (HTTPS or SSL).

<i>Table 8. Communication ports when the target is an IBM Spectrum Protect Plus agent</i>				
Port	Protocol	Initiator	Target	Description
5985	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus appliance	Microsoft Exchange Server	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus appliance	Microsoft Exchange Server	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.

<i>Table 9. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user</i>				
Port	Protocol	Initiator	Target	Description
3260 ¹	TCP	Microsoft Exchange Server	vSnap server	The Microsoft iSCSI Initiator service vSnap target port that is used for mounting LUNS for backup and recovery operations

Table 9. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user (continued)

Port	Protocol	Initiator	Target	Description
137	User Datagram Protocol (UDP)	Microsoft Exchange Server	vSnap server	vSnap Server Message Block (SMB) or Common Internet File System (CIFS) target port that is used for mounting file system shares for transaction log backup and recovery operations
138	UDP	Microsoft Exchange Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations
139	TCP	Microsoft Exchange Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations
443 ²	TCP	Microsoft Exchange Server	vSnap server	Port that allows the agent to communicate with IBM Spectrum Protect Plus for sending alerts in case of log backup failures
445	TCP	Microsoft Exchange Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations

¹ iSCSI initiator is required on this node.

² For Microsoft Exchange Server, this port is available in IBM Spectrum Protect Plus V10.1.4 and later.

Hardware

Table 10. Minimum hardware requirements		
System	Disk Space	Disk Space for Granular Restore Operations
x64: Compatible hardware that is supported by the operating system and Microsoft Exchange Server	A minimum of 500 MB of disk space for the product to be installed	At least 2.1 GB disk space for extra Microsoft prerequisites, which are installed automatically

Db2 requirements

Before you register Db2 with IBM Spectrum Protect Plus, ensure that your system environment meets the outlined requirements.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

The IBM Db2 database backup and restore requirements for IBM Spectrum Protect Plus are as follows.

Configuration requirements

The following IBM Db2 databases are supported:

- IBM Db2 Version 10.5 and later maintenance levels and modification levels: Enterprise Server Edition.
- IBM Db2 Version 11.1 and later maintenance levels and modification levels: Enterprise Server Edition.
- IBM Db2 Version 11.5 and later maintenance levels and modification levels: Enterprise Server Edition.

Operating systems

The following operating systems are supported:

- On PowerPC®:
 - AIX® 7.1 and later modification and fix pack levels (64-bit kernel).
 - AIX 7.2 and later modification and fix pack levels (64-bit kernel).
- On Linux x86_x64:
 - Red Hat Enterprise Linux 6.8 and later maintenance levels and modification levels.
 - Red Hat Enterprise Linux 7 and later maintenance levels and modification levels.
 - SUSE Linux Enterprise Server 11.0 SP4 and later maintenance levels and modification levels.
 - SUSE Linux Enterprise Server 12.0 SP1 and later maintenance levels and modification levels.
 - SUSE Linux Enterprise Server 15.0 SP1 and later maintenance levels and modification levels.
- On Linux on Power® System (little endian)
 - Red Hat Enterprise Linux 7.1 and later maintenance and modification levels.
 - SUSE Linux Enterprise Server 12.0 SP1 and later maintenance and modification levels.
 - SUSE Linux Enterprise Server 15.0 SP1 and later maintenance and modification levels.

Additional notes

Install the latest IBM Db2 patches and updates in your environment.

IBM Db2 pureScale® is not supported

Ensure that your Db2 environment is configured to meet the following criteria:

- Db2 archive logging is activated and Db2 is in recoverable mode.

- Logical volumes holding Db2 table spaces (data and temporary table spaces), the local database directory, and Db2 log files are managed by Logical Volume Manager (LVM2) on Linux and by JFS2 on AIX. LVM2 on Linux and JFS2 on AIX are used for creating temporary volume snapshots. The logical volume grows in size with data as it changes on the source volume while the snapshot exists. For more information, see [“LVM2 and JFS2” on page 191](#).
- Db2 must be in parallel backup mode if multiple partitions are to be protected. Parallel backup mode can be enabled by using Db2 registry variables. For more information, see [“Prerequisites for Db2” on page 189](#).

Software

Review the following software requirements:

- The bash and sudo packages must be installed. Sudo must be version 1.7.6p2 or above. Run `sudo -V` to check the version.
Tip: The required bash and sudo packages are included in the supported Linux86_64 and Linux Power Systems (little endian) operating systems.
- Ensure that the supported version of Linux x86_64, Linux Power Systems (little endian), or AIX is installed.

Connectivity

Ensure that the following connectivity criteria are in place:

- The SSH service is running on port 22 on the server.
- Firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the server using SSH.
- The Secure Shell (SSH) subsystem for SSH is enabled.
- The server can be registered by using a Domain Name System (DNS) name or IP address. DNS names must be resolvable by IBM Spectrum Protect Plus.
- On AIX, ensure that the NFS communication is configured with reserved ports by using the command:
`nfsd -p -o nfs_use_reserved_port=1`.

Authentication and privileges

The Db2 server must be registered in IBM Spectrum Protect Plus by using an operating system user that exists on the Db2 server (referred to as IBM Spectrum Protect Plus agent user).

Ensure that the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.

To use a Db2 database, an IBM Spectrum Protect Plus agent user must have the following permissions:

- Privileges to run commands as the root user and as the Db2 software owner user by using sudo. IBM Spectrum Protect Plus requires these privileges for various tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases.
 - The sudoers configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.
 - The `!requiretty` setting must be set.
- Privileges to read the Db2 inventory using `/usr/local/bin/db2ls`. IBM Spectrum Protect Plus requires this privilege to discover and collect information about Db2 instances and databases.

Ports

The following ports are used by IBM Spectrum Protect Plus agents. Ports that are marked as Accept use a secure connection (HTTPS/SSL).

Table 11. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus virtual appliance. IBM Spectrum Protect Plus virtual appliance contains the following base components: <ul style="list-style-type: none"> • IBM Spectrum Protect Plus server • site • vSnap server • vSnap pool • VADP proxy 	Db2 database	Provides access to troubleshoot and maintain vSnap servers by using the Secure Shell (SSH) protocol

Table 12. Communication ports when the initiator is the IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
111	TCP	Db2 database	vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communications with ONC servers
2049	TCP	Db2 database	vSnap server	Used for Network File System (NFS) data transfer to and from vSnap servers
20048	TCP	Db2 database	vSnap server	Mounts vSnap file systems on clients such as the VMware vStorage API for Data Protection (VADP) proxy, application servers, and virtualization data stores

Hardware

System	Disk Space
Compatible hardware that is supported by the operating system and Db2 database	A minimum of 500 MB of disk space for the product to be installed

MongoDB requirements

Before you register a MongoDB application server with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

Configuration

Database versions:

- MongoDB Version 3.6 and later maintenance and modification levels: Community Server and Enterprise Server Editions.
- MongoDB Version 4.0 and later maintenance levels and modification levels: Community Server and Enterprise Server Editions.

Operating systems:

- On Linux x86_64:
 - Red Hat Enterprise Linux (RHEL) 6.8 and later maintenance and modification levels
 - CentOS 6.8 and later maintenance and modification levels
 - RHEL 7 and later maintenance and modification levels
 - CentOS 7 and later maintenance and modification levels
 - SUSE Linux Enterprise Server 12.0 SP1 and later maintenance and modification levels
- On Linux on Power Systems (little endian):
 - RHEL 7.1 and later maintenance and modification levels
 - CentOS 7 and later maintenance and modification levels

Restriction: On Linux on Power Systems (little endian), only the MongoDB Enterprise Server Edition is supported.

Additional notes

Install the latest MongoDB patches and updates in your environment.

Ensure that your MongoDB environment is configured to meet the following criteria:

- The MongoDB is configured as a stand-alone instance or replica set. Backups of MongoDB sharded cluster instances are not supported. A backup always includes all databases in the instance.
- The MongoDB instance is configured to use the WiredTiger Storage Engine.
- The user in the MongoDB application server registration in IBM Spectrum Protect Plus must be able to retrieve server information and status from the MongoDB admin database.
- Logical volumes of MongoDB data and log paths are managed by Linux Logical Volume Manager (LVM2). LVM2 is used to create temporary volume snapshots. The database files and the journal must be on a single volume. The logical volume grows in size with data as the data changes on the source volume while the snapshot exists. For more information, see [“Linux LVM2 ” on page 254](#).
- In MongoDB, you must use file path names with ASCII characters.

Software

- When the MongoDB application server runs RHEL 6 or CentOS 6, ensure that the `openssl` package is at version 1.0.1e-57 or later. To update the version, run the following command: `yum update openssl`.
- Ensure that the supported version of Linux x86_64 or Linux on Power Systems (little endian) is installed.

Connectivity

Ensure that the following connectivity criteria are in place:

- The Secure Shell (SSH) service is running on port 22 on the server.
- Firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the server by using SSH.
- The secure file transfer protocol (SFTP) subsystem for SSH is enabled.
- The application server can be registered in IBM Spectrum Protect Plus by using a Domain Name System (DNS) name or Internet Protocol (IP) address. DNS names must be resolvable by IBM Spectrum Protect Plus.

Authentication and privileges

- The MongoDB server must be registered with IBM Spectrum Protect Plus by using an operating system user that exists on the MongoDB server (referred to as *IBM Spectrum Protect Plus agent user* for the rest of this topic).
- Ensure that the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- On MongoDB, SSL-based encryption and certificate-based authentication are not supported.
- With the MongoDB Enterprise Server Edition, only Encrypted Storage Engine is supported.

To use a MongoDB database, an IBM Spectrum Protect Plus agent user must have the following permissions:

- Privileges to run commands as root and as a MongoDB software owner user by using sudo. IBM Spectrum Protect Plus requires these privileges for various tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases.
 - The `sudoers` configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.
 - The `!requiretty` setting must be specified.
- Privileges to read the standard MongoDB server module `/usr/local/bin/mongod`. IBM Spectrum Protect Plus requires these privileges to use the PyMongo API to connect to the MongoDB servers by using the instance's assigned DNS name or IP address name and port. This mechanism is used to gather information about MongoDB instances and databases.
- If the MongoDB server is protected by role-based authentication, you must set up the appropriate privileges, see [“Roles for MongoDB” on page 253..](#)

Ports

The following ports are used by IBM Spectrum Protect Plus agent users. The ports use secure connections (HTTPS or SSL).

Table 13. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus virtual appliance ¹	MongoDB	Provides access to troubleshoot and maintain remote proxy host servers running guest applications components by using the Secure Shell (SSH) protocol.

Table 14. Communication ports when the initiator is the IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
111	TCP	MongoDB	vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communications with ONC servers.
2049	TCP	MongoDB	vSnap server	Used for Network File System (NFS) data transfer to and from vSnap servers.
20048	TCP	MongoDB	vSnap server	Mounts vSnap file systems on clients such as the VMware vStorage API for Data Protection (VADP) proxy, application servers, and virtualization data stores.

Table 15. Minimum hardware requirements

System	Disk Space
Compatible hardware that is supported by the operating system and MongoDB.	A minimum of 500 MB of disk space for the product to be installed.

Office 365 requirements

This document details the Microsoft Office 365 backup and restore requirements for IBM Spectrum Protect Plus Version 10.1.5. Before you register a proxy host with IBM Spectrum Protect Plus, ensure that

¹ The IBM Spectrum Protect Plus virtual appliance contains the base components: IBM Spectrum Protect Plus server, site, vSnap server, vSnap pool, and VADP proxy.

the system environment meets the following requirements. The proxy host server is referred to in the user interface (UI) as the *application server*.

Cloud service configuration

To protect a Microsoft Office 365 application, you must register the application with Azure Active Directory and grant the appropriate permissions. To be able to do that, ensure that you have the following items in place:

- Active Microsoft Office 365 subscription.
- Microsoft Office 365 administrative user ID and password.

For more information about how to register Azure, see [Registering with Azure Active Directory](#).

Ensure that you have a Microsoft Office 365 administrative account. Then, you can add users to ensure that they have valid licenses. For instructions, see [Microsoft 365 in Visual Studio subscriptions](#).

Note: The IBM Spectrum Protect Plus server and agent user do not store administrative user IDs or passwords for the Microsoft Office 365 tenant.

Application versions

The following Microsoft Office applications are supported:

- Microsoft Office 365 Business
- Microsoft Office 365 Business Premium
- Microsoft Office 365 Business Essentials
- Microsoft Office 365 Education

Note: The Microsoft Office 365 tenant must be in a global region as defined by Microsoft. National regions are not supported. For more information about regions, see <https://docs.microsoft.com/en-us/graph/deployments>.

Operating systems

The following operating systems are supported:

- On Linux x86_x64
 - Red Hat Enterprise Linux 7.0 and later maintenance and modification levels
 - CentOS 7.0 and later maintenance and modification levels

Additional notes

IBM Spectrum Protect Plus supports the proxy host server running on a physical (bare metal) server and in a virtualization environment.

Software

Ensure that the supported version of Linux x86_64 is installed. In addition, the following software must be installed:

- Java™ 8
- The International Components for Unicode (libicu) rpm-package that corresponds to the installed operating system.

Connectivity

- The Secure Shell (SSH) service must be running on port 22 on the proxy host server. Any firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the proxy host server by using SSH. The secure file transfer protocol (SFTP) subsystem for SSH must also be enabled.

- The server can be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address. DNS names must be resolvable by IBM Spectrum Protect Plus.
- If DNS is not available, you must add the server to the /etc/hosts file on the IBM Spectrum Protect Plus virtual appliance by using the command prompt.

Authentication and privileges

- The agent host must be registered with IBM Spectrum Protect Plus by using an operating system user that exists on the agent host. The agent host user is then referred to as the IBM Spectrum Protect Plus agent user.
- Ensure that the password is correctly configured and that the user can log in without other prompts, such as prompts to reset the password.
- The IBM Spectrum Protect Plus agent user must have privileges to run commands as root using sudo. The sudoers configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.

Network File System (NFS)

The proxy host server must have the native Linux NFS client installed. IBM Spectrum Protect Plus uses NFS to mount storage volumes for backup and restore operations.

Ports

The following ports are used by IBM Spectrum Protect Plus agents users. The ports use secure connections (HTTPS or SSL).

Table 16. Communication ports when the target is an IBM Spectrum Protect Plus agent user.				
Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus virtual appliance. IBM Spectrum Protect Plus virtual appliance contains the following base components: <ul style="list-style-type: none"> • IBM Spectrum Protect Plus server • site • vSnap server • vSnap pool • VADP proxy 	Proxy host server	Provides access to troubleshoot and maintain vSnap servers by using the Secure Shell (SSH) protocol

Table 17. Required communication ports when the initiator is an IBM Spectrum Protect Plus agent user.

Port	Protocol	Initiator	Target	Description
111	TCP	Proxy host server	vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communications with ONC servers
443	TCP	Proxy host server	vSnap server	Allows agents to communicate with IBM Spectrum Protect Plus for sending alerts if log backup failures
2049	TCP	Proxy host server	vSnap server	Used for NFS data transfer to and from vSnap servers
20048	TCP	Proxy host server	vSnap server	Mounts vSnap file systems on clients such as the VMware vStorage API for Data Protection (VADP) proxy, application servers, and virtualization data stores

Hardware

Table 18. Minimum hardware requirements for Office 365

System	Disk Space	Memory
Compatible hardware with quad-core processor supported by the operating system	5 GB of available disk space for temporary files at run time	4 GB Random Access Memory (RAM)

Oracle Server database backup and restore requirements

Review the Oracle database backup and restore requirements for IBM Spectrum Protect Plus.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

Configuration

Database versions

- Oracle 11g R2 Enterprise Edition
- Oracle 12c R1 Enterprise Edition
- Oracle 12c R2 Enterprise Edition
- Oracle 18c Enterprise Edition

- Oracle 19c Enterprise Edition

Note: For multitenant databases in Oracle 12c and later, IBM Spectrum Protect Plus supports protection and recovery of the container database, including all pluggable databases (PDBs) under it. Granular recovery of specific PDBs can be performed by using an Instant Disk Restore recovery combined with the Recovery Manager (RMAN).

Operating systems

On IBM PowerPC:

- AIX 6.1 TL9 and later maintenance and modification levels
- AIX 7.1 and later maintenance and modification levels

On Linux x86_64:

- Red Hat Enterprise Linux (RHEL) 6.5 and later maintenance and modification levels
- RHEL 7.0 and later maintenance and modification levels
- Cent OS 6.5 and later maintenance and modification levels
- Cent OS 7.0 and later maintenance and modification levels
- SUSE Linux Enterprise Server 11.0 SP4 and later maintenance and modification levels
- SUSE Linux Enterprise Server 12.0 SP1 and later maintenance and modification levels
- SUSE Linux Enterprise Server 15.0 and later maintenance and modification levels

Restrictions:

- Oracle DataGuard is not supported.
- Databases must be in ARCHIVELOG mode. IBM Spectrum Protect Plus cannot protect databases running in NOARCHIVELOG mode.
- Real Application Cluster (RAC) database recovery operations are not server pool-aware. IBM Spectrum Protect Plus can recover databases to an RAC, but not to specific server pools.
- RAC databases must be configured such that the RMAN Snapshot Control File location points to shared storage that is accessible to all cluster instances.
- When restoring an Oracle database that was configured for multithreading at the time of backup, the restored database is non-multithreaded. The restored database must be manually reconfigured to use multi-threading.

Software

- The **bash** and **sudo** packages must be installed. The **sudo** package must be version 1.7.6p2 or later. Run **sudo -V** to check the version.

Tip: The required **bash** and **sudo** packages are included in the supported Linux86_64 operating systems.

- **RHEL and CentOS 6 users only:**

To ensure that the `util-linux-ng` package is current, run the following command:

```
yum update util-linux-ng
```

Depending on your version or distribution, the package might be named `util-linux`.

Connectivity

- The Secure Shell (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 by using SSH. The subsystem for SSH must also be enabled.
- The server can be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address. DNS names must be resolvable by IBM Spectrum Protect Plus.

- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus appliance by using the command line.
- When you register Oracle RAC nodes, register each node by using its physical IP or name. Do not use a virtual name or Single Client Access Name (SCAN).

Authentication and privileges

- The Oracle Server must be registered in IBM Spectrum Protect Plus by using an operating system user that exists on the Oracle Server. The user is then referred to as the IBM Spectrum Protect Plus agent user.
- Ensure that the password is correctly configured and that the user can log in without other prompts, such as prompts to reset the password.

To use an Oracle Server, the IBM Spectrum Protect Plus agent user must have the following permissions:

- Privileges to run commands as root and as an Oracle software owner user (for example, `oracle` or `grid`) by using **sudo**. These privileges are required for tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases and Automatic Storage Management (ASM).
 - The `sudoers` configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.
 - The `!requiretty` setting must be set.
 - The `ENV_KEEP` setting must allow the `ORACLE_HOME` and `ORACLE_SID` environment variables to be retained.
- Privileges to read the Oracle inventory. These privileges are required for tasks such as discovering and collecting information about Oracle homes and databases.

To achieve this, the IBM Spectrum Protect Plus agent user must belong to the Oracle inventory group, typically named `oinstall`.

For information about creating a new user with the required privileges, see [“Sample configuration of an IBM Spectrum Protect Plus agent user” on page 47](#).

Network file system (NFS)

The Oracle server must have the native Linux or AIX NFS client installed. IBM Spectrum Protect Plus uses NFS to mount storage volumes for backup and restore operations.

For database restore operations, the Oracle Direct NFS feature is required. IBM Spectrum Protect Plus automatically enables Direct NFS if it is not already enabled.

For Direct NFS to work correctly, the executable `oracle_home/bin/oradism` in each Oracle home directory must be owned by root and have **setuid** privileges. Typically, the binary is preconfigured by the Oracle installer, but on certain systems, this binary might not have the required privileges.

Run the following commands to set the correct privileges:

```
chown root:oinstall ORACLE_HOME/bin/oradism
```

```
chmod 750 ORACLE_HOME/bin/oradism
```

where `oinstall` specifies the group that owns the installation.

Database discovery

IBM Spectrum Protect Plus discovers Oracle installations and databases by searching the `/etc/orainst.loc` and `/etc/oratab` files and the list of running Oracle processes. If the files are not present in their default location, the **locate** utility must be installed on the system so that IBM Spectrum Protect Plus can search for the files.

IBM Spectrum Protect Plus discovers databases and their storage layouts by connecting to running instances and querying the locations of their data files, log files, and other files. In order for IBM Spectrum Protect Plus to correctly discover databases during cataloging and copy operations, databases must be in MOUNTED, READ ONLY, or READ WRITE mode. IBM Spectrum Protect Plus cannot discover or protect database instances that are shut down.

Block change tracking

IBM Spectrum Protect Plus requires Oracle block change tracking to be enabled on protected databases to efficiently perform incremental backups. If block change tracking is not already enabled, IBM Spectrum Protect Plus enables it automatically during the backup job.

To customize the placement of the block change tracking file, you must manually enable the block change tracking feature before you run an associated backup job. If the feature is enabled automatically by IBM Spectrum Protect Plus, the following rules are used to determine the placement of the block change tracking file:

- If the **db_create_file_dest** parameter is set, the block change tracking file is created in the location specified by this parameter.
- If the **db_create_file_dest** parameter is not set, the block change tracking file is created in the same directory as the SYSTEM table space.

Log backup

- The **crond** daemon must be enabled on the application server.
- The IBM Spectrum Protect Plus agent user must have the necessary privileges to use the **crontab** command and create cron jobs. Privileges can be granted through the **crontab.allow** configuration file.

Sample configuration of an IBM Spectrum Protect Plus agent user

The following commands are examples for creating and configuring an operating system user that IBM Spectrum Protect Plus uses to log in to the Oracle Server. The command syntax might vary depending on your operating system type and version.

- Create the user that is designated as the IBM Spectrum Protect Plus agent user:

```
useradd -m sppagent
```

- Set a password:

```
passwd sppagent
```

- If using key-based authentication, place the public key in the `/home/sppagent/.ssh/authorized_keys` directory, or the appropriate file depending on your `sshd` configuration, and ensure that the correct ownership and permissions are set. The commands are structured as shown in the following example:

```
chown -R sppagent:sppagent /home/sppagent/.ssh
chmod 700 /home/sppagent/.ssh
chmod 600 /home/sppagent/.ssh/authorized_keys
```

- Add the user to the Oracle installation and OSDBA group:

```
usermod -a -G oinstall,dba sppagent
```

- If you plan to use ASM, also add the user to the OSASM group:

```
usermod -a -G asmadmin sppagent
```

- Place the following lines at the end of the sudoers configuration file, typically /etc/sudoers. If the existing sudoers file is configured to import a configuration from another directory (for example, /etc/sudoers.d), you can also place the lines in a new file in that directory:

```
Defaults:sppagent !requiretty
Defaults:sppagent env_keep+="ORACLE_HOME"
Defaults:sppagent env_keep+="ORACLE_SID"
sppagent ALL=(ALL) NOPASSWD:ALL
```

Ports

The following ports are used by IBM Spectrum Protect Plus agent users. The ports use secure connections (HTTPS or SSL).

<i>Table 19. Communication ports when the target is an IBM Spectrum Protect Plus agent</i>				
Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus virtual appliance ²	Oracle Server	Provides access to troubleshoot and maintain remote proxy host servers running guest applications components by using the Secure Shell (SSH) protocol

<i>Table 20. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user</i>				
Port	Protocol	Initiator	Target	Description
111	TCP	Oracle Server	vSnap server	Allows Open Network Computing (ONC) clients to discover ports for communications with ONC servers
443	TCP	Oracle Server	vSnap server	Allows agents to communicate with IBM Spectrum Protect Plus for sending alerts if log backups fail.
2049	TCP	Oracle Server	vSnap server	Used for NFS data transfer to and from vSnap servers

² The IBM Spectrum Protect Plus virtual appliance contains the base components: IBM Spectrum Protect Plus server, site, vSnap server, vSnap pool, and VADP proxy.

Table 20. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user (continued)				
Port	Protocol	Initiator	Target	Description
20048	TCP	Oracle Server	vSnap server	Mounts vSnap file systems on clients such as the VMware vStorage API for Data Protection (VADP) proxy, application servers, and virtualization data stores

Hardware

Table 21. Minimum hardware requirements	
System	Disk Space
Compatible hardware that is supported by the operating system and Oracle Server	A minimum of 500 MB of disk space for the product to be installed

Microsoft SQL Server database backup and restore requirements

Review the Microsoft SQL Server database backup and restore requirements for IBM Spectrum Protect Plus.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

Configuration

Database versions

- Microsoft SQL Server 2008 R2 SP3
- Microsoft SQL Server 2012
- Microsoft SQL Server 2012 SP2
- Microsoft SQL Server 2014
- Microsoft SQL Server 2016
- Microsoft SQL Server 2017

Install the latest Microsoft SQL Server patches and updates in your environment.

Operating systems

- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2016
- Microsoft Windows Server 2019

Ensure that Windows Remote Shell (WinRM) is enabled.

An iSCSI route must be enabled between the Microsoft SQL Server system and vSnap server. For more information, see [Microsoft iSCSI Initiator Step-by-Step Guide](#).

IBM Spectrum Protect Plus inventory jobs discover system databases and mark the databases that are eligible for protection. Log backups are marked as ineligible for all system databases and databases running in simple recovery model.

Microsoft SQL Server backup and restore operations require that the Windows PowerShell binary path is set in the %PATH% environment variable.

Microsoft SQL Server backup operations are limited to less than 64 TB volume size. If you plan to back up test mode restored databases, use the global preference to limit the size of backup target volumes to a size less than 64 TB. You must set this global preference before you run the first backup for the service level agreement (SLA) that protects the databases. If the size of the backup target volumes is 64 TB or more, the backup fails.

In-Memory online transaction processing (OLTP)

In-Memory online transaction processing (OLTP) is a memory-optimized database engine that is used to improve database application performance. This engine is supported in Microsoft SQL Server 2014 and later. The following requirements and limitations apply to In-Memory OLTP usage:

- The restore file path is limited to 256 or fewer characters. If the original path exceeds this length, consider using a customized restore file path to reduce the length.
- The metadata that can be restored is subject to Volume Shadow Copy Service (VSS) and Microsoft SQL Server restore capabilities.

Incremental backups

IBM Spectrum Protect Plus uses update sequence number (USN) change journal technology to perform incremental backups in a Microsoft SQL Server environment. The USN change journal provides write range tracking for a volume when the file size meets the minimum file size threshold requirement. The changed bytes offset and length extent information can be queried against a specific file.

To enable write range tracking, the system environment must meet the following requirements:

- Windows Server 2012 R2 or later
- New Technology File System (NTFS) Version 3.0 or later

The following technologies are not supported for changed bytes tracking:

- Resilient File System (ReFS)
- Server Message Block (SMB) 3.0 protocol
- SMB TFO (Transparent Failover)
- SMB 3.0 with Scale-Out file shares (SO)

By default, 512 MB of space is allocated for USN change journaling. In addition, when journal overflow is detected, the allocated space doubles in size when an overflow is detected, to a maximum of 2 GB.

The minimum space required for shadow copy storage is 100 MB, although more space might be required on systems with increased activity. If the free space on the source volume is less than 100 MB, the Microsoft SQL Server agent checks the source volume space and causes a backup to fail. A warning message is displayed in the job log when free space is less than 10%, and then the backup proceeds.

A base backup is forced when the following conditions are detected:

- Journal discontinuity is reported. This condition can occur when the log reaches the maximum size, when journaling is disabled, or when the cataloged USN ID is changed.
- The file size is less than or equal to the tracking threshold size, which by default is 1 MB.
- A file is added after a previous backup job.

Log backups

IBM Spectrum Protect Plus supports log backups:

With staging area

IBM Spectrum Protect Plus uses the backup folder that is configured for the Microsoft SQL Server instance to stage the collection of logs, before copying log files to the vSnap repository. Sufficient free

space must be available to store transaction logs in a log backup. The staging area can be modified by changing the backup folder configuration by using SQL Server Management Studio (SSMS).

Without staging area

This type of log backup requires vSnap Active Directory (AD) integration. To learn how to configure the vSnap server, see [Setting vSnap Active Directory](#).

The Microsoft SQL service user must be in an Active Directory (AD) domain. The SLA policy must be configured to use a site that contains a vSnap that is integrated to the same AD domain as the SQL Server service user.

To ensure that SQL Server Log Backup works properly, a Windows Group Policy change might be required. The Group Policy Object (GPO) setting for the **Network security: LAN Manager authentication level policy** at **Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options**, must be set to one of the following options:

- **Not Defined**
- **Send NTLMv2 response only.**
- **Send NTLMv2 response only. Refuse LM.**
- **Send NTLMv2 response only. Refuse LM & NTLM.**

The **Send NTLM response only** option is not compatible with the vSnap Common Internet File System (CIFS) and SMB version and can cause CIFS authentication problems.

Configuring always on availability groups

Configure the preferred instance for backup operations by using Microsoft SQL Server Management Studio. Complete the following steps:

1. Select the **Availability Group** node.
2. Select the availability group you that you want to configure. Then, select **Properties**.
3. In the **Availability Group Properties** dialog box, select **Backup Preferences**.
4. In the **Where should backups occur** pane, select any option.

When a secondary replica is preferred, and more than one secondary replica is available, the IBM Spectrum Protect Plus job executor selects the first secondary replica in the preferred list reported by the IBM Spectrum Protect Plus SQL Server agent.

The Microsoft SQL Server agent sets the VSS backup type to COPY_ONLY.

The **No Recovery** option does not support production mode restore operations for SQL AlwaysOn availability groups.

Registration and authentication

Register each Microsoft SQL server with IBM Spectrum Protect Plus by name or IP address. When you register an SQL Server Cluster node, register each node by name or IP address.

Restriction: The IP address must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. Both servers must have a Windows Remote Management service that is listening on port 5985. The fully qualified domain name must be resolvable and can be routed from the IBM Spectrum Protect Plus server and from the vSnap server.

The user identity must have sufficient rights to install and start the IBM Spectrum Protect Plus Tools Service on the node. These rights include **Log on as a service** rights. For more information, see the Microsoft article: [Add the Log on as a service Right to an Account](#)

If the virtual machine 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.

Kerberos

Kerberos-based authentication can be enabled by specifying a configuration file on the IBM Spectrum Protect Plus appliance. The settings override the default Windows NT LAN Manager (NTLM) protocol.

For Kerberos-based authentication only, the user identity must be specified in the `username@FQDN` format. The user name must be able to authenticate by using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain specified by the fully qualified domain name.

Privileges

To use a Microsoft SQL Server, an IBM Spectrum Protect Plus agent user must have the following permissions:

- Microsoft SQL Server `public` and `sysadmin` permissions
- Windows local administration permission, which are required by the VSS framework, and volume and disk access
- Permissions to access cluster resources in an SQL Server AlwaysOn and SQL Server FCI environment.

Every Microsoft SQL Server instance can use a specific user account to access the resources of that SQL Server instance.

The SQL Server Virtual Device Interface (VDI)-based framework is used to interact with SQL Server databases and to log backup and restore operations. A VDI connection requires Microsoft SQL Server `sysadmin` permissions. The owner of a restored database is not changed to the original owner. A manual step is required to modify the owner of a restored database. For more information about the VDI framework, see the Microsoft article: [SQL Server VDI backup and restore operations require Sysadmin privileges](#)

The target Microsoft SQL Server service account must have permissions to access SQL Server restore files. See *Administrative Considerations* in the Microsoft article: [Securing Data and Log Files](#)

The Windows Task Scheduler is used to schedule log backups. Depending on the environment, users might receive the following error:

A specified logon session does not exist. It might already have been terminated.

This behavior occurs when a network access Group Policy setting is enabled. For instructions about disabling the setting, see the Microsoft Support article: [A specified logon session does not exist. It may already have been terminated, error when you try to map to a network drive of a DFS share](#)

Ports

The following ports are used by IBM Spectrum Protect Plus agent users. The ports use secure connections (HTTPS or SSL).

Table 22. Communication ports when the target is an IBM Spectrum Protect Plus agent				
Port	Protocol	Initiator	Target	Description
5985	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus virtual appliance ³	Microsoft SQL Server	Provides access to the Microsoft Windows Remote Management (WinRm) service for Windows-based servers.

³ The IBM Spectrum Protect Plus virtual appliance contains the base components: IBM Spectrum Protect Plus server, site, vSnap server, vSnap pool, and VADP proxy.

Table 22. Communication ports when the target is an IBM Spectrum Protect Plus agent (continued)

Port	Protocol	Initiator	Target	Description
5986	TCP	IBM Spectrum Protect Plus virtual appliance	Microsoft SQL Server	Provides access to the Microsoft Windows Remote Management (WinRm) service for Windows-based servers.

Table 23. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user

Port	Protocol	Initiator	Target	Description
3260 ⁴	TCP	Microsoft SQL Server	vSnap server	The Microsoft iSCSI Initiator service vSnap target port that is used for mounting LUNS for backup and recovery operations
137	User Datagram Protocol (UDP)	Microsoft SQL Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations
138	UDP	Microsoft SQL Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations
139	TCP	Microsoft SQL Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations

⁴ iSCSI initiator is required on this node.

Table 23. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user (continued)

Port	Protocol	Initiator	Target	Description
443 ⁵	TCP	Microsoft SQL Server	vSnap server	Port that allows the agent to communicate with IBM Spectrum Protect Plus for sending alerts in case of log backup failures
445	TCP	Microsoft SQL Server	vSnap server	vSnap SMB or CIFS target port that is used for mounting file system shares for transaction log backup and recovery operations

Hardware

Table 24. Minimum hardware requirements

System	Disk Space
Compatible hardware that is supported by the operating system and Microsoft SQL Server	A minimum of 500 MB of disk space for the product to be installed

Kubernetes Backup Support requirements

Before you deploy IBM Spectrum Protect Plus Kubernetes Backup Support in the Kubernetes environment, ensure that your system environment meets the outlined requirements.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see [technote 2013790](#).

Kubernetes Backup Support is available only in English in IBM Spectrum Protect Plus Version 10.1.5.

Container versions

Docker containers are supported in Kubernetes Backup Support.

Operating systems

On Linux x86_64:

- Red Hat Enterprise Linux (RHEL) 7.6
- RHEL 7.7

Additional requirements

- Kubernetes 1.13 and later patches and updates
- Kubernetes 1.14 and later patches and updates
- Kubernetes 1.15 and later patches and updates

⁵ For Microsoft SQL Server, this port is available in IBM Spectrum Protect Plus V10.1.4 and later.

- Kubernetes 1.16 and later patches and updates
- Ceph Container Storage Interface (CSI) driver 1.1 with Rados Block Device (RBD) storage

To install and configure container backup support, the backup administrator must deploy the Kubernetes Backup Support software in the Kubernetes environment. For instructions, see [“Installing Kubernetes Backup Support”](#) on page 313.

Software

- Kubernetes Backup Support protects only persistent storage that was allocated by a storage plug-in that supports the Container Storage Interface (CSI).
- Only formatted volumes can be mounted to the data mover for copy operations.
- Ensure that Kubernetes Metrics Server 0.3.5 or later is installed and running on your cluster. The metrics server is required for the Kubernetes Backup Support scheduler to determine the resources that are used for multiple concurrent data mover instances. For more information, see [“Verifying whether the metrics server is running”](#) on page 314.
- Copy backup and snapshot restore operations require the **VolumeSnapshotDataSource** alpha feature to be enabled. To enable the **VolumeSnapshotDataSource** alpha feature, you must patch the Kubernetes scheduler, controller, and API server. For instructions, see [“Enabling the VolumeSnapshotDataSource feature”](#) on page 313.
- Ensure that the following cluster prerequisites are met:
 - You must be running a Kubernetes cluster with CSI support.
 - Persistent storage must be provided by the CSI driver, which must support CSI snapshot capabilities.
 - A storage class must be defined for the persistent volumes that are being protected.
 - The Kubernetes command-line tool **kubectl** must be accessible on the installation host and in the local path.
 - CSI snapshot support must be enabled on the **kubectl** command line.
 - The target image registry must be accessible from the Kubernetes cluster. The target image registry can be a local image registry or an external image registry. For an external image registry, you can configure the image pull secret to secure your environment.
 - The Kubernetes Backup Support product installation package must be on the master node or another administration node. The administration node must have similar access to the master node with regards to Docker, the **kubectl** tool, and the cluster image registry.
 - To create new cluster-wide resources, you must be logged in to the target cluster as a user with **cluster-admin** privileges.
 - Ensure that Kubernetes Backup Support secrets that include user IDs, passwords, and keys are encrypted at rest in the etcd distributed key-value store. For more information, see [Encrypting Secret Data at Rest](#).

Helm prerequisites

The Helm tool must be configured on the target cluster so that a new deployment can be run with the **helm** command line. Deploying a package with Helm enables cluster-wide role-based access control (RBAC) rules and role bindings to be generated.

For the Kubernetes cluster, to install Helm as root user with the Kubernetes administrative user account, run the following script, which is included in the installation package:

```
./helm_install_k8s.sh
```

IBM Spectrum Protect Plus prerequisites

External, non-container components such as IBM Spectrum Protect Plus and the IBM Spectrum Protect Plus vSnap server must be provisioned and configured by the IBM Spectrum Protect Plus administrator.

- An administrative account for Kubernetes Backup Support must be configured on IBM Spectrum Protect Plus.

This administrative account can be configured as a global Lightweight Directory Access Protocol (LDAP) account in the data center. This global account is required for access to all external components that Kubernetes Backup Support operates with.

You must specify this account name in the `BAAS_ADMIN` parameter in the `baas_config.cfg` configuration file before you deploy Kubernetes Backup Support. The `baas_config.cfg` is located in the `installer` directory. For instructions, see [“Installing and deploying Kubernetes Backup Support images”](#) on page 315.

- An IBM Spectrum Protect Plus instance must be deployed and licensed as a VMware virtual appliance.
Network connectivity must exist to and from the target cluster. The IBM Spectrum Protect Plus IP address and port number must be specified in the `baas_config.cfg` file before you deploy Kubernetes Backup Support. Only one port (443) can be specified for use with all IBM Spectrum Protect Plus instances.
- An IBM Spectrum Protect Plus vSnap instance must be deployed as a VMware virtual appliance.
 - Network connectivity must exist to and from the target Kubernetes cluster and IBM Spectrum Protect Plus vSnap instance.
 - The vSnap instance must be configured as an external vSnap server for storing backups. For instructions, see Chapter 3, [“Installing vSnap servers,”](#) on page 67.
 - If backups are encrypted at rest, ensure that enough capacity is allocated for encryption on the vSnap server.

Connectivity

Ensure that the following connectivity criteria are in place:

- SSH service is running on Kubernetes NodePort services.
- Firewalls must be configured to allow IBM Spectrum Protect Plus to connect data mover containers by using SSH over the NodePort port range of the Kubernetes cluster. The NodePort service allows the specific port in the NodePort range to be determined by Kubernetes at run time.
- The server can be registered in IBM Spectrum Protect Plus by using a Domain Name System (DNS) name or an Internet Protocol (IP) address. DNS names must be resolvable by IBM Spectrum Protect Plus.

Authentication and privileges

Ensure that you specify the user name for the IBM Spectrum Protect Plus administrative account and data mover in the `baas_config.cfg` configuration file. For more information, see [“Installing and deploying Kubernetes Backup Support images”](#) on page 315.

To access the device that is associated with the persistent volume, the data mover container must be a privileged container.

Ports

The following communications ports are used by IBM Spectrum Protect Plus agents. The ports use secure connections (HTTPS or SSL).

Table 25. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
Assigned by the NodePort service in Kubernetes	TCP	IBM Spectrum Protect Plus virtual appliance ⁶	Kubernetes	Used by IBM Spectrum Protect Plus to connect to the data mover container to deploy and run agents.

For SSH connections between containers in the Kubernetes environment, port 22 is used. For everywhere else, whether on the Kubernetes hosts or outside the cluster, the port that the NodePort service assigned at runtime is used.

Table 26. Communication ports when the initiator is the IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
111	TCP	Kubernetes	vSnap server	Allows ONC clients to discover ports for communications with ONC servers
443	TCP	Kubernetes	vSnap server	Used for IBM Spectrum Protect Plus issued commands to run backup, restore, inventory, and other configuration operations
2049	TCP	Kubernetes	vSnap server	Used for NFS data transfer to and from vSnap servers
20048	TCP	Kubernetes	vSnap server	Mounts vSnap file systems on clients such as the VADP proxy, application servers, and virtualization data stores

Related concepts

[“Protecting containers” on page 307](#)

⁶ Refers to the IBM Spectrum Protect Plus server, which is a component of the IBM Spectrum Protect Plus virtual appliance.

Kubernetes Backup Support is a feature of IBM Spectrum Protect Plus that extends data protection to containers in Kubernetes clusters. Kubernetes is a system for orchestrating containers across clusters of hosts.

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 a 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 1072392](#).

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 103](#)

You can update the IBM Spectrum Protect Plus virtual appliance, vSnap servers, and the VADP proxy servers to get the latest features and enhancements. Software patches and updates are installed by using the IBM Spectrum Protect Plus administrative console or command-line interface for these components.

Related tasks

[“Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 58](#)

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 60](#)

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. A local vSnap server that is already named and registered is also installed on the virtual appliance.

[“Installing a vSnap server” on page 67](#)

When you deploy an IBM Spectrum Protect Plus appliance, a vSnap server is automatically installed. This server is the primary backup destination. In larger enterprise environments, additional vSnap servers might be 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

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in [“Component requirements” on page 11](#) and [“Hypervisor requirements” on page 26](#).

- Download the virtual appliance template installation file CC1QCML . ova from Passport Advantage® Online. For information about downloading files, see [technote 1072392](#).
- 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 62](#).

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 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 appliances.
- IBM Spectrum Protect Plus has not been tested for IPv6 environments.

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 CC1QCML . 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 compute resource. Click **Next**.
5. Review the template details. Click **Next**.

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.

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. For vSphere Web Client, enter the property values for the virtual appliance: 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: For vSphere Client, 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](#).

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, the IBM Spectrum Protect Plus application as well as a local vSnap server which is built into it will be registered and installed on it automatically. To start IBM Spectrum Protect Plus, 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 .
Upload the product key.	See “Uploading the product key” on page 62 .
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 91 .

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. A local vSnap server that is already named and registered is also installed on the virtual appliance.

Before you begin

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in [“Component requirements” on page 11](#) and [“Hypervisor requirements” on page 26](#).
- Download the installation file CC1QDML.exe from Passport Advantage Online. For information about downloading files, see [technote 1072392](#).
- 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 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.

Procedure

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

1. Copy the CC1QDML .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 alliances 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 pane, 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 alliance. 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 62](#).

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 62 .
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 91 .

Assigning a static IP address

To reassign a new static IP address after initial deployment, a network administrator can assign a 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 91](#).
2. From a CentOS 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 58](#)

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 60](#)

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. A local vSnap server that is already named and registered is also installed on the virtual appliance.

Uploading the product key

IBM Spectrum Protect Plus runs in an evaluation mode for a limited time period. A valid product key is required to enable IBM Spectrum Protect Plus features indefinitely.

Before you begin

Save the product key to a computer with internet access and record the location of the key.

Procedure

To upload the product key, 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, select **Authentication Type > System**. Enter the `serveradmin` password to access the Administration Console. 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 91](#).

3. Click **Manage your licenses**.
4. Click **Choose File**, and then browse for the product key on your computer,
5. Click **Upload new license**.
6. Click **Logout**.

What to do next

After you upload the product key, 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 91 .

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.

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

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

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
```

Installing iSCSI initiator utilities

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

About this task

iSCSI initiator utilities can be installed on the IBM Spectrum Protect Plus appliance or a vSnap server. The iSCSI initiator utilities are delivered along with IBM Spectrum Protect Plus, but are not installed automatically. To install the utilities, follow the procedure.

Procedure

1. Log on to the appliance or server that is to be directly connected to the iSCSI mounted storage.
 - For the IBM Spectrum Protect Plus appliance, use the Secure Shell (SSH) protocol and authenticate with the appropriate administrative credentials.
 - For a vSnap server, 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
```

Chapter 3. Installing vSnap servers

Every installation of IBM Spectrum Protect Plus requires at least one vSnap server, which is the primary backup destination.

In both VMware and Hyper-V environments, one vSnap server with the name localhost is automatically installed when the IBM Spectrum Protect Plus appliance is initially deployed. An onboard vSnap server resides on a partition of the IBM Spectrum Protect Plus appliance and is registered and initialized in IBM Spectrum Protect Plus. In smaller backup environments, the onboard vSnap server might be sufficient.

In larger enterprise environments, 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 appliance is installed and deployed. After installation, 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

When you deploy an IBM Spectrum Protect Plus appliance, a vSnap server is automatically installed. This server is the primary backup destination. In larger enterprise environments, additional vSnap servers might be required.

Before you begin

Complete the following steps:

1. Review the vSnap system requirements in [“Component requirements”](#) on page 11.
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, see [technote 1072392](#).

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 requirements”](#) on page 18 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.
3. Run `setenforce 0` to apply the setting immediately without requiring a restart.
4. Download the vSnap installation file CC1QGML .run from Passport Advantage Online. For information about downloading files, see [technote 1072392](#).

5. Make the file executable through the command `chmod +x CC1QGML.run`, and then run the executable. The vSnap packages are installed, plus all of required components.

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 Chapter 4, “Managing vSnap servers,” on page 73.

Installing a virtual vSnap server and a VADP proxy in a VMware environment

To install a virtual vSnap server and a vStorage API for Data Protection (VADP) proxy in a VMware environment, deploy an Open Virtualization Format (OVF) template. This creates a machine that contains the vSnap server and the VADP proxy.

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 62. Work with your network administrator when configuring network properties.

Procedure

1. Download the server and proxy template installation file `CC1QEML.ova` from Passport Advantage Online. For information about downloading files, see [technote 1072392](#).
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 `CC1QEML.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 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 a disk format to store the virtual disks. To optimize performance, you can select thick provisioning, which is preselected. Thin provisioning requires less disk space, but might impact performance. Click **Next**.
10. 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**.
11. Provide details of the VADP configuration, including the IP address of the IBM Spectrum Protect Plus appliance.

For ESXi server 5.5, this prompt is shown when the OVF deployment template reaches the **Properties** step.

For the ESXi server 6.0 and later, this prompt is shown when the OVF deployment template reaches the **Customize Template** step.

12. 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.

13. Click **Next**.

14. Review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template. Deployment might take significant time.

15. After the OVF template is deployed, power on your newly created virtual machine. You can power on the VM from the vSphere Client.

Important: The VM must remain powered on for the IBM Spectrum Protect Plus application to be accessible.

16. 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 Chapter 4, “Managing vSnap servers,” on page 73.
Configure the VADP environment.	See “Setting options for VADP proxies” on page 163.

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 CC1QFML . exe from Passport Advantage Online. For information about downloading files, see [technote 1072392](#).
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 62.

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 Chapter 4, “Managing vSnap servers,” on page 73.

Uninstalling a vSnap server

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

Before you begin

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 > Hypervisors > VMware**.

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 91.

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 84.

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.

Chapter 4. Managing vSnap servers

To enable backup and restore jobs, at least one IBM Spectrum Protect Plus virtual appliance and at least one vSnap server is required. The vSnap server can be located on the IBM Spectrum Protect Plus appliance or on its own appliance, or it can be a physical vSnap installation. Each vSnap server location must be added so that IBM Spectrum Protect Plus recognizes it.

Adding a vSnap server as a backup storage provider

The onboard vSnap server is registered in IBM Spectrum Protect Plus when the appliance is deployed. You must add any additional servers that are installed on either virtual or physical appliances so that they are recognized by IBM Spectrum Protect Plus.

Before you begin

After you add a vSnap server as a backup storage provider, you might have to configure and administer certain aspects of vSnap, such as network configuration or storage pool management. For more information, see [“vSnap server administration reference”](#) on page 84.

Procedure

To add 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 91.
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 pane, click **System Configuration > Backup Storage > Disk**.
5. Click **Add Disk Storage**.
6. Complete the fields in the **Storage Properties** pane:

Hostname/IP

Enter the resolvable IP address or hostname 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.

Username

Enter the user name for the vSnap server that you created in step [“2”](#) on page 73.

Password

Enter the password for the user.

7. Click **Save**.

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 80.
Expand the vSnap storage pool.	See “Expanding a vSnap storage pool” on page 81.
If necessary, configure and administer certain aspects of vSnap, such as network configuration or storage pool management.	“vSnap server administration reference ” on page 84

Related tasks

[“Start IBM Spectrum Protect Plus”](#) on page 91


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 pane, click **System Configuration > Backup Storage > Disk**.
2. Click the edit icon  that is associated with a vSnap server.
The **Edit Storage** pane is displayed.
3. Revise the vSnap server settings, and then click **Save**.

Deleting a vSnap server

You can delete a vSnap server that is no longer used in your IBM Spectrum Protect Plus environment.

Before you begin

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



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

Before you delete a vSnap server, review the scenarios to determine whether deletion 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 delete the vSnap server. If you delete 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 delete the vSnap server. If you delete 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 74.

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

- Delete 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. Delete 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 [“Adding a vSnap server as a backup storage provider” on page 73](#).
 - 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 delete 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 [“Replacing a failed vSnap server” on page 83](#).

Procedure

To delete a vSnap server, complete the following steps:



1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. Click the delete icon  that is associated with a vSnap server.
3. Click **Yes** to delete the server from IBM Spectrum Protect Plus.

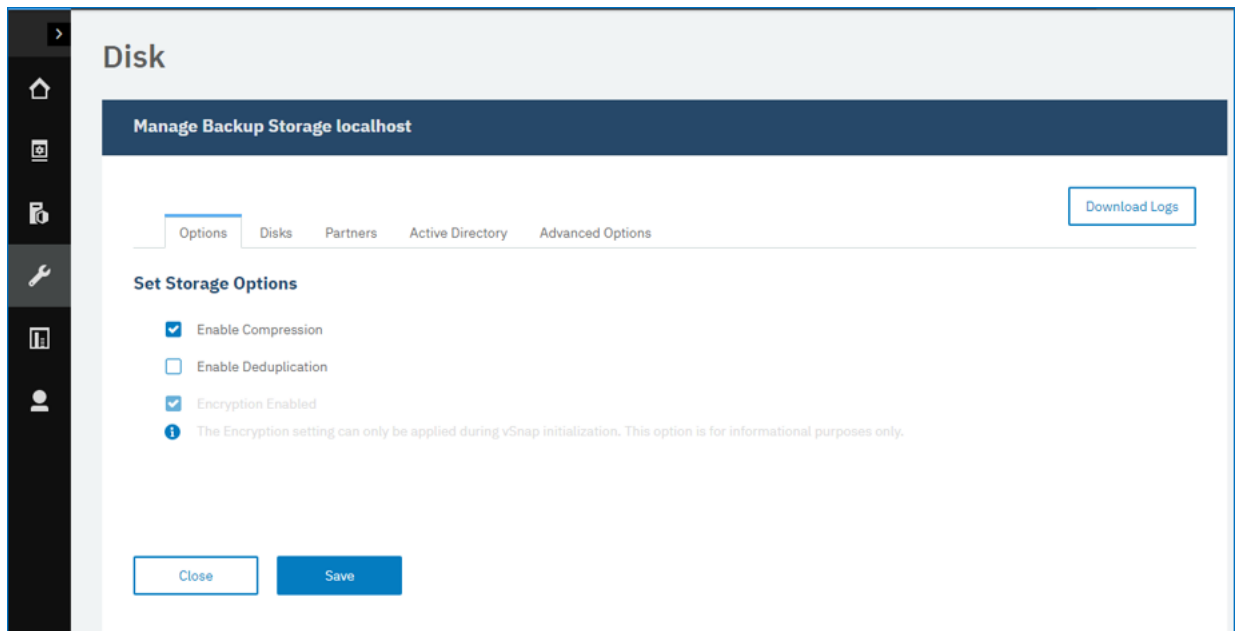
Configuring backup storage options

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

Procedure

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

1. In the navigation pane, click **System Configuration  Backup Storage > Disk**.
The **Disk Storage** table lists the hostname of primary and secondary sites with the version and the capacity usage.
2. In the **Disk Storage** pane, click the manage icon  that is associated with the disk that you want to update.
3. Select from the storage options as shown.



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.



4. Click **Save**.

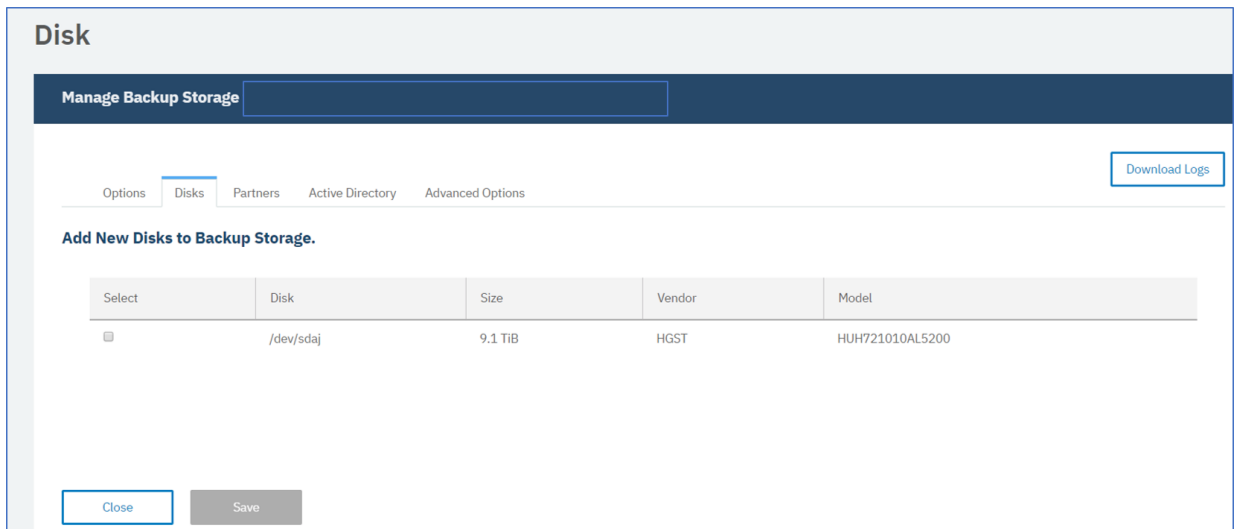
Adding new disks to backup storage

If you require more space for backup operations in a selected storage pool, you can add unused disk storage. This applies to primary and secondary backup storage.

Procedure

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

1. In the navigation, click **System Configuration** , **Backup Storage** > **Disk**.
2. In the **Disk Storage** pane, click the manage icon  that is associated with the server that you want to edit.
3. Select a disk to add to your storage environment from the list of available disks in the **Add New Disks to Backup Storage** table.




4. Click **Save**.

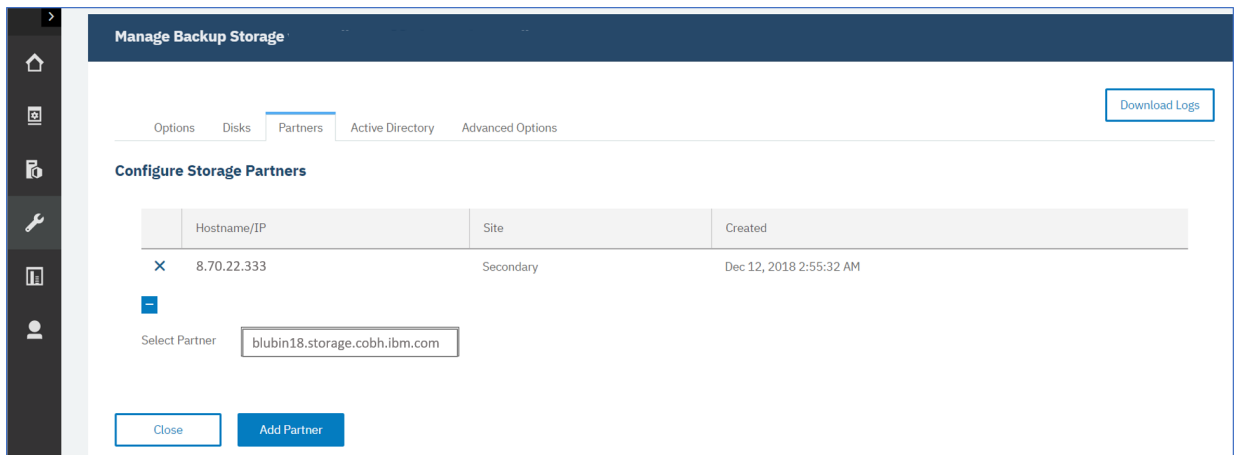
Configuring backup storage partners

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.

Procedure

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

1. In the navigation, click **System Configuration** , **Backup Storage** > **Disk**.
Configured partners that are already added are listed in the table.
2. In the **Partners** pane, select a partner to add to you primary or secondary backup storage host from the drop-down menu.



3. Click **Add Partner** to add the partner and close the window.

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.

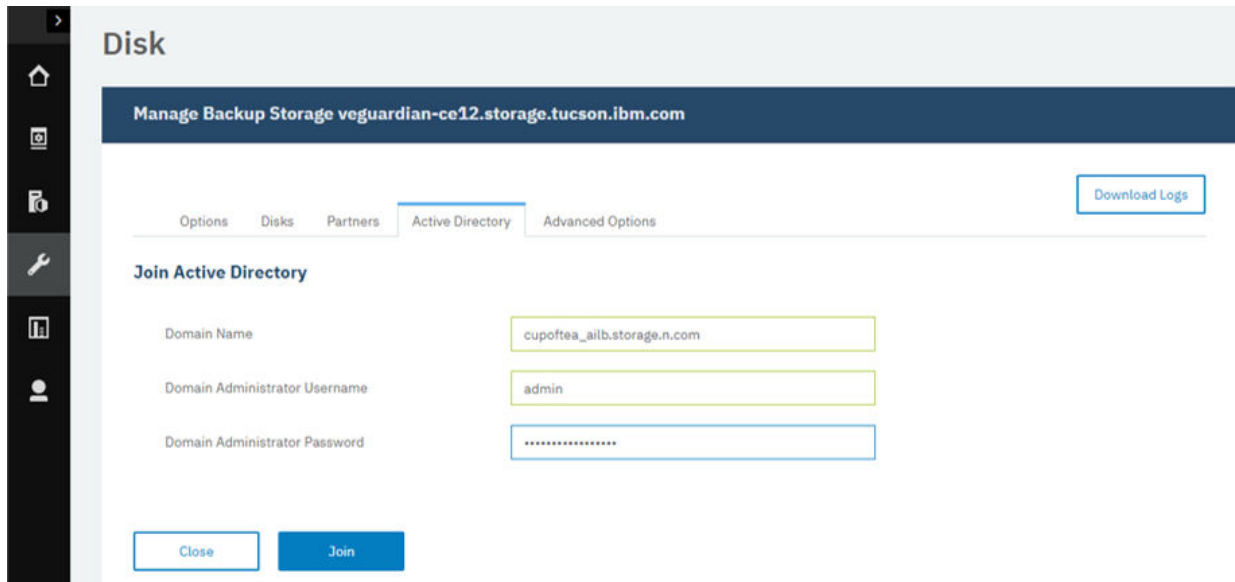
Before you begin

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 pane, click **System Configuration** , **Backup Storage** > **Disk**.
2. On the **Active Directory** tab, click the manage icon  that is associated with the primary or secondary host that you want to edit.
3. Enter the domain name of the Active Directory, along with the user name and password for the Active Directory adminstraro as shown in the following picture.



The screenshot shows the 'Disk' management interface. At the top, there's a header 'Disk' and a sub-header 'Manage Backup Storage veguardian-ce12.storage.tucson.ibm.com'. Below this, there are tabs: 'Options', 'Disks', 'Partners', 'Active Directory' (selected), and 'Advanced Options'. A 'Download Logs' button is in the top right. The 'Join Active Directory' section contains three input fields: 'Domain Name' with the value 'cupoftea_aillb.storage.n.com', 'Domain Administrator Username' with the value 'admin', and 'Domain Administrator Password' with masked characters. At the bottom, there are 'Close' and 'Join' buttons.



4. 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 pane, click **System Configuration** , **Backup Storage** > **Disk**.
2. In the **Manage Backup Storage** pane, click the manage icon  that is associated with the host that you are managing.
3. On the **Advanced Options** tab, configure advanced options as shown in the following example:

Disk

Manage Backup Storage localhost

Options Disks Partners Active Directory **Advanced Options**

Set Advanced Options

Concurrent stream limit for copy to archive object storage

Concurrent stream limit for copy to standard object storage

Concurrent stream limit for replication


Rate limit per stream in bytes/second for replication and copy to standard object storage

Retrieval tier for restore from AWS archive object storage (Bulk, Standard, or Expedited)

ConcurrentBackup

Figure 5. 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.
- **Rate limit per stream in bytes/second for replication and copy to standard object storage:** This value defines the maximum transfer rate in bytes per second that the backup host uses for each data stream when you are replicating or copying data to standard Object Storage. The specified value is the maximum in the absence of any other limiting factors. The actual rate of each data stream can be less than this value and depends on available system resources, network conditions, and any bandwidth throttling defined in site options.
- **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 backups can run.
 - Pause: to pause the backup job until the previous job completes.
 - Limit: to set a maximum limit on the number of backup jobs that can run concurrently..

- 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**.
4. Click **Close**.

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 you must run only for new installations.

About this task

As part of the initialization process, vSnap creates a storage pool using any available unused disks on the system. The OVA-based deployments of vSnap each contain a default 100 GB unused virtual disk which is used to create the pool.

If no unused disks are found, the initialization process completes without creating a pool.

For information about how to expand, create, and administer storage pools, see [“Storage management” on page 84](#).

You can use the IBM Spectrum Protect Plus user interface or the vSnap server console to initialize vSnap servers.

For servers that are deployed in a virtual environment, the user interface provides a simple method to run the initialization operation.

For servers that are deployed in a physical environment, 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.

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.

About this task

For the onboard vSnap installation that is registered as part of an IBM Spectrum Protect Plus installation, you are prompted to start the initialization process the first time you log in to the user interface. No further steps are required.

Procedure

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

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. From the **Actions** menu that is associated with the server, select the initialization method:

Initialize with Encryption

Enable encryption of backup data on the vSnap server.

Initialize

Initialize the vSnap server without encryption enabled.

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 a physical 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`. 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 91.

You can also use a user ID that has vSnap admin 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 84.

2. Run the **vsnap system init --skip_pool** command. The command requires no further interaction and completes all initialization tasks except for the creation of a storage pool. The process might take 5 - 10 minutes to complete.

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 84.

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, either by adding virtual disks to the vSnap virtual machine or adding physical disks to the vSnap physical server. See the vSphere documentation for information about creating additional virtual disks.

Before you begin

Virtual or physical disks must be added to the vSnap server prior to this procedure. Expanding existing volumes is not supported.

Procedure

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

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. Select **Actions > Rescan** for the vSnap server that you want to rescan.
3. Click the manage icon  that is associated with the vSnap server, and then expand the **Add New Disks to Backup Storage** section.
4. Add and save the selected disks. The vSnap pool expands by the size of the disks that are added.

Establishing a replication partnership for a vSnap server



By using backup storage replication, you can asynchronously backup data from one vSnap server to another.

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 establish a replication partnership, complete the following steps:

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. Click the manage icon  that is associated with the vSnap server that you want to add a replication partnership to, and then expand the **Configure Storage Partners** section.
3. Click the add icon .
4. From the **Select Partner** list, select a vSnap server with which to establish a replication partnership.
5. Click **Add Partner**.

What to do next


After you create a replication partnership, complete the following action to enable replication:

Action	How to
Select the Backup Storage Replication option in the SLA policy that is associated with the backup job.	See “Creating an SLA policy” on page 145

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 pane, click **System Configuration > Site** to open the **Site Properties** pane.
2. Click the edit icon  that is associated with the site for which you want to change the throughput.
3. Click **Enable Throttle**.

The rate of the throughput is displayed in MB/s.

4. Adjust the throughput:
 - Change the rate of throughput with the up and down arrows.
 - Change the data value. The choices include Bytes/s, KB/s, MB/s, or GB/s.

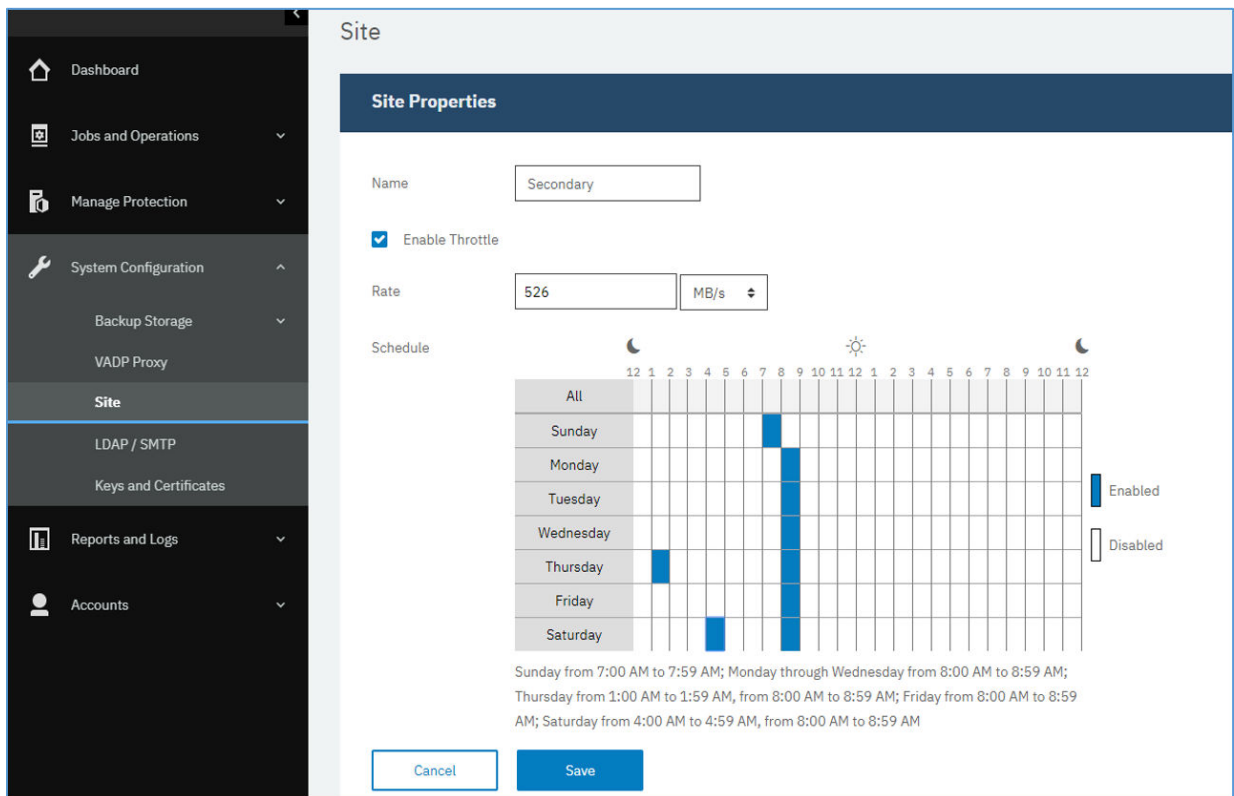


Figure 6. Enabling different throttles for different times to improve throughput

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

Note: To clear a timeslot, click the timeslot. 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](#).

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 command-line interface is the primary means of administering vSnap. 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 admin privileges. Use the **vsnap user create** command to create additional operating system users that have these 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 91](#).

The command line interface consists of several commands and sub-commands that manage various aspects of the system. See [“Storage management” on page 84](#) and [“Network management” on page 87](#) for details on using these commands. 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 of the most common operations can also be completed from the IBM Spectrum Protect Plus user interface. Log in to the user interface and click **System Configuration > Backup Storage > Disk** in the navigation pane. Click the manage icon  for a vSnap server to edit its settings.

Related tasks

[“Managing vSnap servers” on page 73](#)

To enable backup and restore jobs, at least one IBM Spectrum Protect Plus virtual appliance and at least one vSnap server is required. The vSnap server can be located on the IBM Spectrum Protect Plus appliance or on its own appliance, or it can be a physical vSnap installation. Each vSnap server location must be added so that IBM Spectrum Protect Plus recognizes it.

[“Configuring advanced storage options” on page 78](#)

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

Storage management

You can configure and administer storage pools for a vSnap server.

Managing disks

vSnap creates a storage pool using disks 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 SAN storage attached to the physical server. The local disks may already have external redundancy enabled via a hardware RAID controller, but if not, vSnap can also create RAID-based storage pools for internal redundancy.

Disks that are attached to vSnap servers must be thick provisioned. If disks are thin provisioned, the vSnap server will not have an accurate view of free space in the storage pool, which might lead to data corruption if the underlying datastore runs out of space.

If vSnap was deployed as part of a virtual appliance, it already contains a 100 GB starter virtual disk that can be used to create a pool. 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. Alternatively, creating new SLA policies will force backups to use an alternate vSnap.

It is essential to protect against corruption caused by a VMware datastore on a vSnap server reaching its capacity. Create a stable environment for virtual vSnap servers that do not use RAID configurations by utilizing thick provisioned VMDKs. Replicating to external vSnap servers provides further protection.

A vSnap server will become invalidated if the vSnap pool is deleted or if a vSnap disk is deleted in a non-redundant RAID configuration. All data on the vSnap server will be lost. If your vSnap server becomes invalidated you must unregister the vSnap server using the IBM Spectrum Protect Plus interface, then run the maintenance job. Once complete, the vSnap server can be re-registered.

Managing 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 generated upon pool creation are stored under the directory `/etc/vsnap/keys/` on each vSnap server. For disaster recovery purposes, back up the key files manually outside the vSnap server. After a pool is created, use the following commands as the `serveradmin` user to copy them to a temporary location and then copy them to a desired, secure backup location outside the vSnap host.

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

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

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: Click **System Configuration > Backup Storage > Disk** in the navigation pane, and then click the **Actions** menu next to the relevant vSnap server and select **Rescan**.

Showing disks

Run the **`vsnap disk show`** command to list all disks that are on the vSnap system,

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, otherwise they are marked as used by the partition table or file system that is discovered on them.

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

Showing storage pool information

Run the **vsnap pool show** command to view information about each storage pool.

Creating a storage pool

If you completed the simple initialization procedure described in [“Completing a simple initialization”](#) on page 80, 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 85. For information about available options, pass the **--help** flag for any command or subcommand.

Specify a user-friendly 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 choose to enable compression and deduplication for the pool during creation. You can also update the compression/deduplication settings at a later time by using the **vsnap pool update** command.

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

raid0

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

Once a disk has been added to a raid0 pool it cannot be removed. Disconnecting the disk will result in the pool becoming unavailable, which can be resolved only by destroying and recreating the pool.

raid5

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

raid6


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

Expanding a storage pool

Before expanding a pool, ensure that one or more unused disks are available as described in [“Showing disks”](#) on page 85.

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, pass the **--help** flag for any command or subcommand.

User interface: Click **System Configuration > Backup Storage > Disk** in the navigation pane. Click the manage icon  for a vSnap server to manage it, and then expand the **Add New Disks** tab. The tab displays all unused disks discovered on the system. Select one or more disks and click **Save** to add them to the storage pool.

Network management

Configure and administer network services for a vSnap server.

Showing network interface information

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

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

mgmt

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

nfs

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

iscsi

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

smb

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

repl

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

Modifying services associated with network interfaces

Run the **vsnap network 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.

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**, **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.

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 vSnap is installed or updated, Linux kernel Version 4.19.xxxx is installed by default. If you opt out of the kernel upgrade to V4.19.xxxx and remain on the V3.10.xxxx, a kernel V3.10.xxxx that is compatible with the vSnap server is installed and used. In both cases, kernel headers and tools associated with the kernel are not installed. If you plan to compile or use customer 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 91.

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:

```
$ 4.19.xxxx
```

3. After the kernel version is determined, navigate to the appropriate directory.

For kernel version 4.19.xxxx, navigate to this directory:

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

For kernel version 3.10.xxxx, 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. To install the kernel header or tool, issue the following command:

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

Results

The kernel header or tool is installed.

Chapter 5. Getting off to a quick start

To start using IBM Spectrum Protect Plus, you must complete steps that include defining resources that you want to protect and creating service level agreement (SLA) policies, also known as backup policies, for those resources. This getting started section provides the basic steps 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.

Before you start, 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 “[Product deployment roadmap](#)” on [page 11](#) 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, hypervisor and database application 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 admin user account with the password password. The administrator is prompted to reset the user name for this account after logging in. The administrator cannot reset the user name to admin, root, or test.</p> <p>After the initial startup, the application administrator can start the application by using this user account or another account that the infrastructure administrator creates.</p>

Action	Owner	Description
“Manage sites” on page 92	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 and Secondary, but a custom site can also be created and assigned when the vSnap server is added.</p> <p>Before continuing to the following actions, review the available sites and determine whether you want to add new sites or modify the existing ones.</p>
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 servers for hypervisors or database applications that host data that you want to protect.
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.

Action	Owner	Description
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 a number of 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://host_name
```

Where *host_name* is the IP address of the virtual machine where the application is deployed. This connects you to IBM Spectrum Protect Plus.

2. Enter your user name and password to log on.

If this is your first time logging on, the default user name is `admin` and the password is `password`. You are prompted to reset the default user name and password. You cannot reset the user name to `admin`, `root`, or `test`.

3. Click **Sign In**.

4. If you are logging on to IBM Spectrum Protect Plus for the first time, you are prompted to complete the following actions:

- 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.
- Start the initialization process for the onboard vSnap server. Select **Initialize** or **Initialize with encryption enabled** to encrypt data on the server.

Manage sites

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.

About this task

A site is assigned to a vSnap server when the server is added to IBM Spectrum Protect Plus. Review the available sites by clicking **System Configuration > Site** in the navigation pane 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 and Secondary sites.

The Demo site is available only for the onboard vSnap server. You cannot use this site with any other vSnap server.

Procedure

To add or edit a site, complete the following steps:

1. In the navigation pane, click **System Configuration > Site**.
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 Enable Throttle to manage the throughput for site replication and copy operations as described in “Adding a site” on page 124.d. Click Save.
Edit a site.	<ol style="list-style-type: none">a. Click Edit Site.b. Click the edit icon  that is associated with a site.c. Optional: Select Enable Throttle to manage the throughput for site replication and copy operations as described in “Editing a site” on page 125.d. Click Save.

Related concepts

[“Product components” on page 1](#)

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

[“Managing sites” on page 124](#)

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 include the frequency and retention of backups.

About this task

The three default SLA policies are Gold, Silver, and Bronze. You can use these policies as they are or modify the policies. You can also create custom SLA policies.

If a virtual machine 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.

For example purposes, this task does not include instructions for enabling replication for vSnap servers or for copying data to secondary backup storage, which are optional features. For information about how to set up these features in the SLA policy, see [“Creating an SLA policy” on page 145](#).

Backup copies of data are called snapshots.

Procedure

To create an SLA policy, complete the following steps:

1. In the navigation pane, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy**.
The **New SLA Policy** pane is displayed.
3. In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
4. In the **Operational Protection** section under **Main Policy**, set the following options for backup operations. These operations occur on the vSnap servers that are defined in the **System Configuration > Backup Storage > Disk** window.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this check box to create the main policy without defining a frequency or start time. Policies created without a schedule can be run on-demand.

Frequency

Enter the frequency for backup operations.

Start Time

Enter the date and time that you want the backup operation to start.

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 check box 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 no encrypted vSnap servers are available, the associated job will fail.

The following example shows a new SLA policy named Copper that runs every 3 days at midnight with a retention of 1 month:

Policy Overview

New SLA Policy

Name: Copper

Operational Protection

Main Policy

Retention: 1 Months

☐ Disable Schedule

Frequency: 3 Days

Start Time: 01/29/2019 00:00

Target Site: Primary

☐ Only use encrypted disk storage

Replication Policy

☐ Backup Storage Replication

☐ Disable Schedule

Frequency: 1 Days

Start Time: 01/29/2019 01:00

Target Site: Secondary

☐ Only use encrypted disk storage

☒ Same retention as source selection

Cancel Save

Figure 7. Creating an SLA policy

5. Click **Save**. The SLA policy can now be applied to backup job definitions as shown in [“Add resources to a job definition”](#) on page 98.

Related concepts

[“Replicate backup-storage data ”](#) on page 5

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.

[“Copy snapshots to secondary backup storage”](#) on page 6

The vSnap server is the primary backup location for snapshots. All IBM Spectrum Protect Plus environments have at least one vSnap server. Optionally, you can copy snapshots from a vSnap server to secondary backup storage.

[“Managing SLA policies for backup operations”](#) on page 145

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 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 hypervisors or applications that are 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. 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 371](#).

To create custom user roles and resource groups, see [“Creating a resource group” on page 364](#) and [“Creating a role” on page 368](#)

Procedure

To create an account for an application administrator, complete the following steps:

1. In the navigation pane, 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 name and password for the application administrator.
5. In the **Assign Role** section, select **VM Admin**.

The permissions are shown in the **Permission Groups** section.

The screenshot shows the 'Add User' pane in the VMware vSphere interface. The pane is titled 'User' and 'Add User - User Information and Role'. It contains a form for creating a new user. The 'Select the type of user or group you want to add.' dropdown is set to 'Individual new user'. The 'Username' field contains 'vmadmin' with a note 'Username must not be 'root', 'admin' or 'test''. The 'Password' field is masked with dots and has a 'Show' button, with a note 'Password must contain at least 8 characters.' Below the form is the 'ASSIGN ROLE' section with checkboxes for 'Application Admin', 'Backup Only', 'Restore Only', 'SYSADMIN', 'Self Service', and 'VM Admin' (which is checked). At the bottom is the 'PERMISSION GROUPS' section with expandable buttons for 'Certificate' and 'Cloud'. 'Cancel' and 'Continue >' buttons are at the bottom of the pane.

Figure 8. Creating a user account and assigning a role

6. Click **Continue**.

7. In the **Add Users - Assign Resources** section, select the **All Resources** resource group, and then click **Add resources**.

The resource group is added to the **Selected Resources** section.

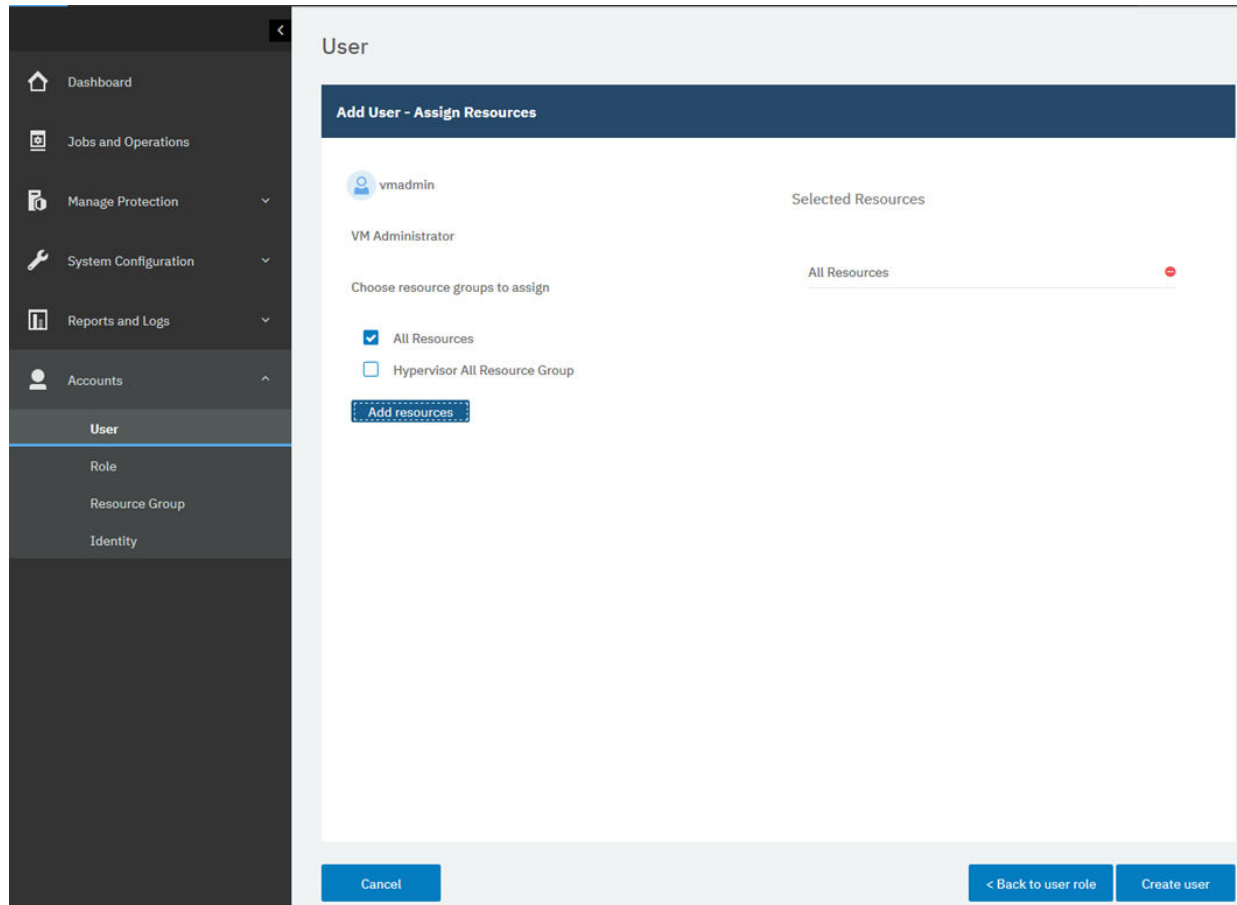


Figure 9. Selecting a resource group for the user account

8. Click **Create user**.

Related concepts

[“Managing user access” on page 363](#)

By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Add resources to protect

Resources are servers for hypervisors or applications that host data 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 resource. To add other resources, see the instructions by resource type in [Chapter 9, “Protecting hypervisors,” on page 151](#) and [Chapter 10, “Protecting applications,” on page 189](#).

Procedure

To add a vCenter Server instance, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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 SSL** check box to enable an encrypted Secure Sockets Layer (SSL) connection. The typical default port is 80 for non SSL connections or 443 for SSL 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 following example shows populated fields.

The screenshot displays the VMware Backup application interface. On the left is a dark sidebar with a navigation menu. The main content area is light gray and titled 'VMware'. At the top of this area is a 'Manage vCenter' button with a gear icon and a 'Create job' button. Below these is a dark blue header for the 'Manage vCenter' section. The main part of the page contains two sections: 'vCenter Properties' and 'Options'. The 'vCenter Properties' section includes input fields for 'Hostname/IP' (containing '192.0.2.0'), 'Use existing user' (an unchecked checkbox), 'Username' (containing 'admin_192.0.2.0'), 'Password' (masked with dots), 'Port' (containing '443'), and a checked 'Use SSL' checkbox. The 'Options' section contains a field for 'Maximum number of VM's to process concurrently per ESX server' with the value '3'. At the bottom of the form are 'Cancel' and 'Save' buttons. A dark blue footer bar at the very bottom is labeled 'VMware Backup'.

Figure 10. 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 possible 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 select an SLA policy for resources that are in a VMware vCenter. To select a policy for other resources, see the instructions by resource type in [Chapter 9, “Protecting hypervisors,”](#) on page 151 and [Chapter 10, “Protecting applications,”](#) on page 189.

Procedure

To select an SLA policy, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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.

The following example shows a specific hard disk that is selected for backup:

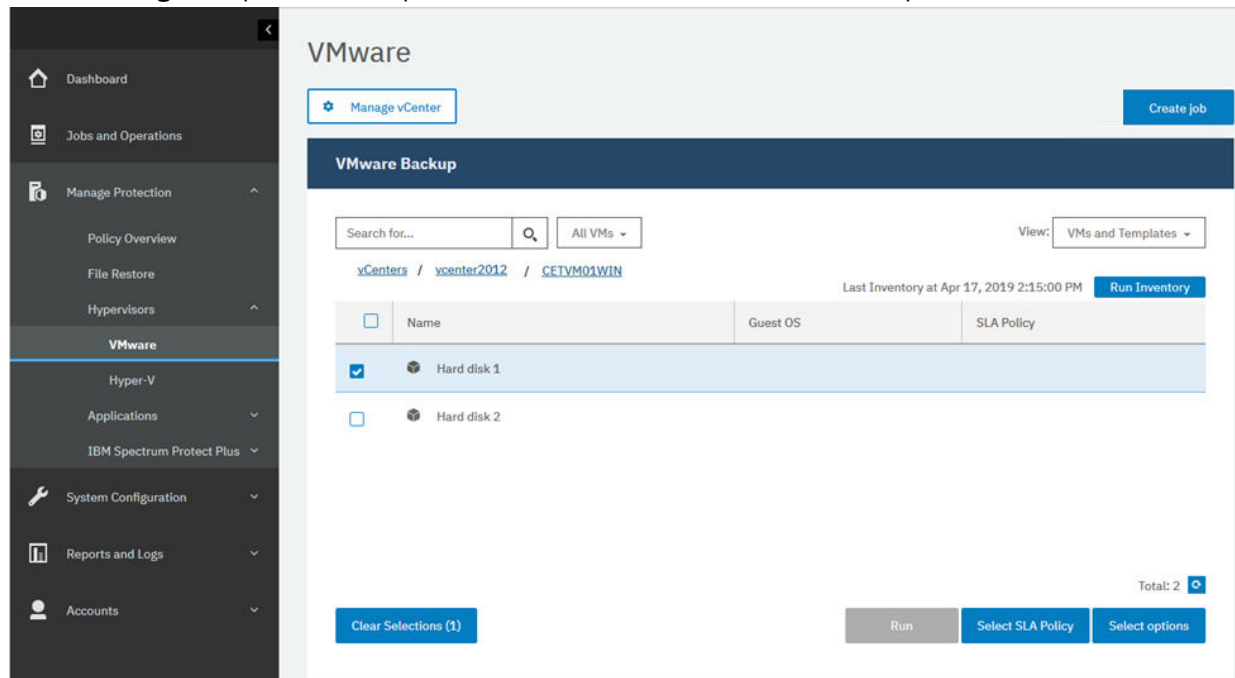


Figure 11. Selecting resources for backup

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup data criteria to the job definition.

The following example shows the SLA policy **Copper** selected:

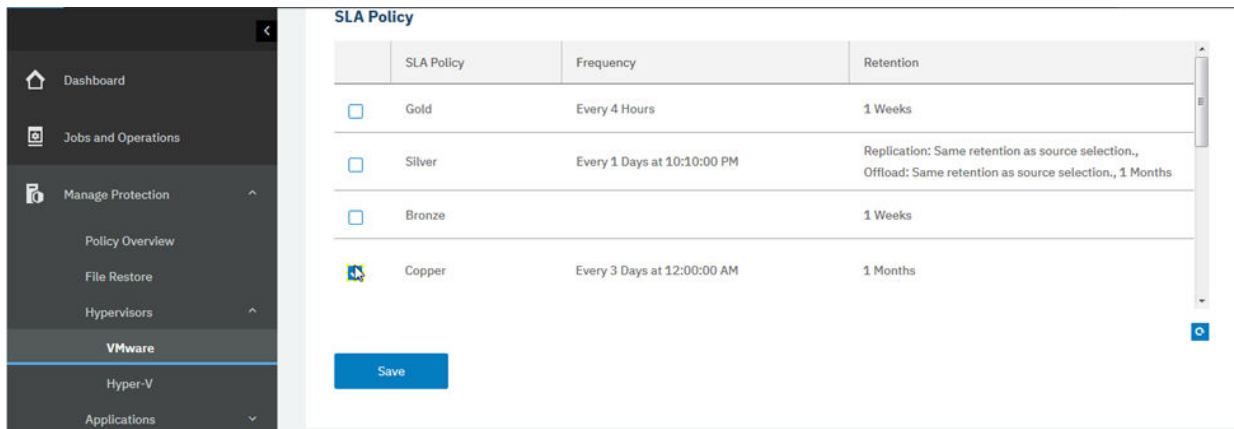


Figure 12. Selecting an SLA policy

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 155.
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 159.

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** and then clicking the **Policy and Job List** tab. For instructions, see [“Start a backup job”](#) on page 99.

Related concepts

[“Protecting IBM Spectrum Protect Plus”](#) on page 339

Protect the IBM Spectrum Protect Plus application by backing up the underlying databases for disaster recovery scenarios. Configuration settings, registered resources, restore points, backup storage settings, and job information are backed up to a vSnap server that is defined in the associated SLA policy.

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, 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 **Actions** > **Start** as shown in the following example:

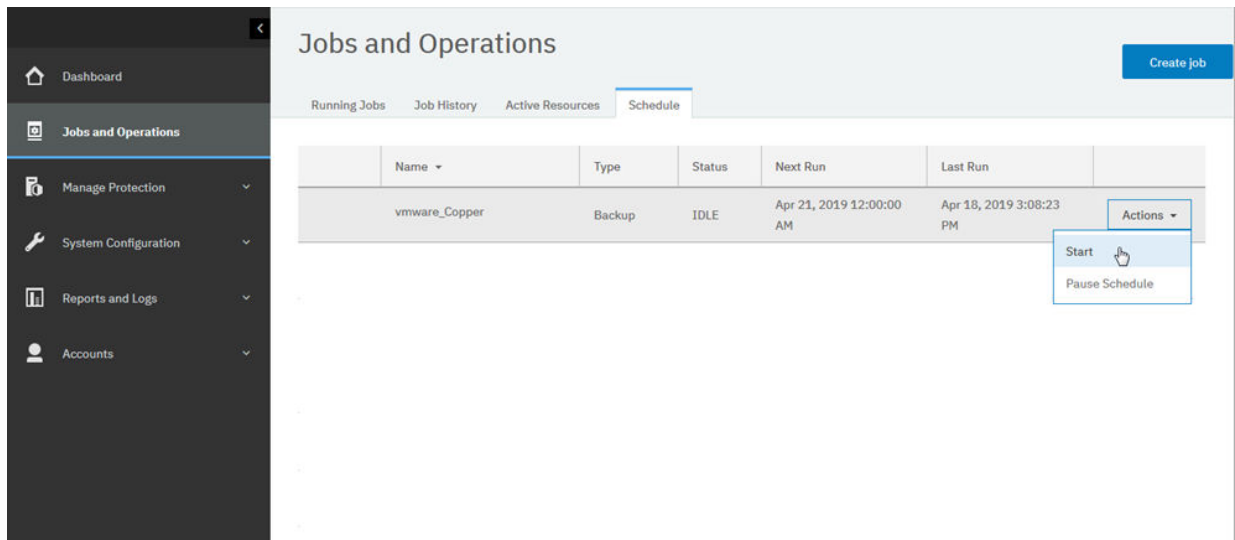


Figure 13. Starting a job

3. To view the job log in detail, 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 an information 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.

4. You can download a job log from the page by clicking **Download .zip**. If you want to cancel the job, click **Actions > Cancel**.

5. Click the **Actions** menu that is associated with the job that you want to start and click **Start** as shown in the following example:

Related concepts

[“Managing jobs and operations” on page 343](#)

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 predefined default parameters or custom parameters.

Procedure

To run a report, complete the following steps:

1. In the navigation pane, click **Reports and Logs > Reports**.
2. Expand a report type and select a report to run as shown in the following example:




Figure 14. Selecting a report to run

3. Run the report either with custom parameters or default parameters:

- To run the report with custom parameters, set the parameters in the **Options** section, 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 353](#)

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 6. Updating IBM Spectrum Protect Plus components

You can update the IBM Spectrum Protect Plus virtual appliance, vSnap servers, and the VADP proxy servers to get the latest features and enhancements. Software patches and updates are installed by using the IBM Spectrum Protect Plus administrative console or command-line interface for these components.

For information about available update files and how to obtain them from an IBM download site, see [technote 1072392](#).

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:

- You must separately update vSnap servers that are not on IBM Spectrum Protect Plus virtual appliances.
- The update process through the administrative console 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 administrative console is the primary means for updating IBM Spectrum Protect Plus features and underlying infrastructure components including the operating system and file system.

Take the following actions:

- 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 application”](#) on page 339.
- After IBM Spectrum Protect Plus is updated, it cannot roll back to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of your environment before you update 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.

Updating the IBM Spectrum Protect Plus virtual appliance

Use the IBM Spectrum Protect Plus administrative console to update the virtual appliance. Updating IBM Spectrum Protect Plus can be run offline or online if you have external internet access.

Before you begin

You can update IBM Spectrum Protect Plus V10.1.2 or later directly to the current version. If you are using V10.1.1, you must update to V10.1.2 and then update to the current version. For instructions about how to update from V10.1.1 to V10.1.2, see [Updating the IBM Spectrum Protect Plus virtual appliance to version 10.1.2](#).

Before you begin the update process, complete the following steps:

1. 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 application”](#) on page 339.

2. For offline updates, download the prerequisite IBM Spectrum Protect Plus update named CC1QHML.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 during the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED.

For a list of download images, including the required operating system update for the virtual appliance, see [technote 1072392](#).

About this task

When you have access to the internet, you can choose to run the update procedure online. If you do not have access to the internet, you can run the offline update procedure.

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 in as an IBM Spectrum Protect Plus user with SYSADMIN privileges, enter your administrator user name and password. If you log in by using the admin user account, you are prompted to reset the user name and password. You cannot reset the user name to admin, root, or test.
System (recommended)	To log on as a system user, enter the serveradmin password. The default password is sppDP758-SysXyz. 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**.

Note: You can select only one update file at a time.

When the update completes, the virtual machine where the application is deployed automatically restarts.

Important: After the IBM Spectrum Protect Plus update completes, you must update any external vSnap and VADP proxy servers in your environment.

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 pane, click **Jobs and Operations**, and then click the **Schedule** tab.

Find the jobs that you paused.

9. From the **Actions** menu for the paused jobs, select **Release Schedule**.

Related tasks

[“Updating vSnap servers” on page 105](#)

The default vSnap server is updated with the IBM Spectrum Protect Plus appliance. You must update additional vSnap servers that are installed on either virtual or physical appliances separately.

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 V10.1.5 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 vSnap servers

The default vSnap server is updated with the IBM Spectrum Protect Plus appliance. You must update additional vSnap servers that are installed on either virtual or physical appliances separately.

Before you begin

You can update your vSnap servers directly from version 10.1.2 or later to the current version. If you are using version 10.1.1, you must update to version 10.1.2 and then update to the current version. For instructions about how to update to version 10.1.2, see [Updating vSnap servers to version 10.1.2](#).

Test restore jobs need to complete prior to initiating an update to vSnap. Jobs that are not completed or canceled when an upgrade 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 11.](#)

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 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 Enterprise Linux documentation.

Related tasks

[“Updating a vSnap server” on page 107](#)

The default vSnap server is updated with the IBM Spectrum Protect Plus appliance. You must update additional vSnap servers that are installed on either virtual or physical appliances separately.

Updating the operating system for a virtual vSnap server

Updating the vSnap server operating system provides 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.

Before you begin

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 application ” on page 339.](#)
2. If you are updating from IBM Spectrum Protect Plus 10.1.1, you must update to version 10.1.2 and then update to the current version. For instructions about how to update to version 10.1.2, see [Updating vSnap servers to version 10.1.2.](#) For more information in determining the version of the vSnap server and operating system, see [“Updating vSnap servers” on page 105.](#)
3. For information on obtaining the ISO file, see [“Updating the IBM Spectrum Protect Plus virtual appliance” on page 103.](#)

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 `CC1QHML.iso`. Move the ISO file to the `/tmp` directory on the vSnap server and rename the file to `spp_with_os.iso`.

```
$mv CC1QHML.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. Proceed with the instructions found in the [“Updating a vSnap server” on page 107](#) topic. When the CC1QGML.run file is executed, the installer will optionally update the operating system if /tmp/spp_with_os.iso is present.

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 107](#)

The default vSnap server is updated with the IBM Spectrum Protect Plus appliance. You must update additional vSnap servers that are installed on either virtual or physical appliances separately.

Updating a vSnap server

The default vSnap server is updated with the IBM Spectrum Protect Plus appliance. You must update additional vSnap servers that are installed on either virtual or physical appliances separately.

Before you begin

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 application ” on page 339](#).
2. If you are updating from IBM Spectrum Protect Plus 10.1.1, you must update to version 10.1.2 and then update to the current version. For instructions about how to update to version 10.1.2, see [Updating vSnap servers to version 10.1.2](#).
3. Download the vSnap update file CC1QGML.run and copy it to a temporary location on the vSnap server. For information about downloading files, see [technote 1072392](#).

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 CC1QGML.run file is located, make the file executable and run the installer by issuing the following commands:

```
chmod +x CC1QGML.run
```

```
sudo ./CC1QGML.run
```

The vSnap packages are installed.

3. Start the updated version of IBM Spectrum Protect Plus.
4. In the navigation pane, click **Jobs and Operations**, and then click the **Schedule** tab.
Find the jobs that you paused.
5. From the **Actions** menu for the paused jobs, select **Release Schedule**.

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 application ” on page 339](#).

Note: If the VADP proxy is not registered with IBM Spectrum Protect Plus, the VADP component that is packaged in the vSnap appliance will not be updated. Only VADP proxies registered with IBM Spectrum Protect Plus will 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 options icon  and select **Set Options**. Through the User setting, select a previously entered username and password 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.

If you are attempting to update as a non-root user, special instructions will need to be followed in order to push-install or push-update a VADP proxy.

1. Create a file in the /etc/sudoers.d/ directory.

```
sudo cd /etc/sudoers.d/
```

2. Write the text to the file and save it by pressing CTRL+D on the keyboard when done.

```
sudo cat > 99-vadpuser
Defaults !requiretty
vadpuser ALL=NOPASSWD: /tmp/cdm_guestapps_vadpuser/runcommand.sh
<<Press CTRL+D>>
```

3. Set the appropriate permissions on the file.

```
sudo chmod 0440 99-vadpuser
```

What to do next

After you update the VADP proxies, complete the following action:

Action	How to
Run the VMware backup job.	<p>See “Backing up VMware data” on page 155.</p> <p>The proxies are indicated in the job log by a log message similar to the following text:</p> <p>Run remote vmdkbackup of MicroService: http://<proxy nodename, IP:proxy_IP_address</p>

Related tasks

[“Creating VADP proxies”](#) on page 161

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 63

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 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 Secure Sockets Layer (SSL) certificates.

In most cases, IBM Spectrum Protect Plus is installed on a virtual appliance. The virtual appliance contains the application and the inventory. Maintenance tasks are completed in vSphere Client, 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 ESX infrastructure, or a user with an understanding of IBM Spectrum Protect Plus, VMware, and Linux command-line usage.

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



Attention: Update underlying components of IBM Spectrum Protect Plus only by using the update facilities that are provided by IBM.

Managing secondary backup storage

The vSnap server is the primary backup location for snapshots. All IBM Spectrum Protect Plus environments have at least one vSnap server. Optionally, you can copy snapshots from a vSnap server to a cloud storage system or a repository server.

For information about copying snapshot data to secondary storage, see [“Copy snapshots to secondary backup storage”](#) on page 6.

Managing cloud storage

You can copy snapshot data to cloud storage for longer-term data protection.

Adding Amazon S3 Object Storage

You can add Amazon Simple Storage Service (S3) as a backup storage provider to IBM Spectrum Protect Plus to enable copy operations to Amazon S3 storage.

Before you begin

Configure the key that is required for the cloud object. For instructions, see [“Adding an access key”](#) on page 121.

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 menu, click **System Configuration > Backup Storage > Object Storage**.
2. Click **Add Object Storage**.
3. From the **Provider** list, select **Amazon S3**.
4. Complete the fields in the **Object Storage Registration** form:

Name

Enter a meaningful name that helps you to identify the cloud storage.

Region

Select the Amazon Web Services (AWS) regional endpoint of the cloud storage.

Use existing 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 help 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.

Enable Deep Archive

Optionally select this option to enable the Amazon S3 Glacier Deep Archive storage class.

5. Click **Get Buckets** to connect IBM Spectrum Protect Plus to AWS to retrieve the list of available buckets.

6. Select the bucket that you plan to use as the copy target.

The **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.

7. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.

8. Optional: In the **Archive object storage bucket** field, 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 [“Copy snapshots to secondary backup storage” on page 6](#).

9. Select **Deep Archive** to register Amazon S3 Glacier Deep Archive Buckets for long-term archiving.

10. Click **Register** to complete the operation.

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” on page 145 . To modify an existing SLA policy, see “Editing an SLA policy” on page 149 .

Adding IBM Cloud Object Storage as a backup storage provider

Add IBM Cloud Object Storage to enable IBM Spectrum Protect Plus to copy data to IBM Cloud.

Before you begin

Configure the key and certificate that are required for the cloud object. For instructions, see [“Adding an access key” on page 121](#) and [“Adding a certificate” on page 122](#).

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 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 menu, click **System Configuration > Backup Storage > Object Storage**.
2. Click **Add Object Storage**.
3. From the **Provider** list, select **IBM Cloud Object Storage**.
4. Complete the fields in the **Object Storage Registration** pane:

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Select the endpoint of the cloud storage.

Use existing 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 help 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, then click **Upload**.

Copy and paste

Select to enter the name of the certificate, copy and paste the contents of the certificate, 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. Click **Get Buckets**, and then select a bucket to serve as the copy target.

After the buckets are generated, the **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.

6. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.
7. Optional: In the **Archive object storage bucket** field, 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 [“Copy snapshots to secondary backup storage” on page 6](#).
8. Click **Register**.
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.	<p>To create an SLA policy, see “Creating an SLA policy” on page 145.</p> <p>To modify an existing SLA policy, see “Editing an SLA policy” on page 149.</p>

Adding Microsoft Azure cloud storage as a backup storage provider

Add Microsoft Azure cloud storage to enable IBM Spectrum Protect Plus to copy data to Microsoft Azure Blob storage.

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 how to create buckets, see Azure documentation.

Procedure

To add Microsoft Azure cloud storage as backup storage provider, complete the following steps:

1. In the navigation pane, click **System Configuration > Backup Storage > Object Storage**.
2. Click **Add Object Storage**.
3. From the **Provider** list, select **Microsoft Azure Blob Storage**.
4. Complete the fields in the **Object Storage Registration** pane:

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Select the endpoint of the cloud storage.

Use existing 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 help identify the key.

Storage Account Name

Enter the Microsoft Azure access storage account name. This is 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, either key1 or key2.

5. Click **Get Buckets**, and then select a bucket to serve as the copy target.

After the buckets are generated, the **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.

6. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.
7. Optional: In the **Archive object storage bucket** field, 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 [“Copy snapshots to secondary backup storage”](#) on page 6.

8. Click **Register**.

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.	To create an SLA policy, see “Creating an SLA policy” on page 145. To modify an existing SLA policy, see “Editing an SLA policy” on page 149.

Adding S3 compatible object storage

In addition to backing up data to Amazon Simple Storage Service (S3) and IBM Cloud Object Storage, you might want to back up data to other S3 compatible object storage providers. Before you back up data in a production environment to any other S3 compatible object storage, ensure that the object storage has been validated for use with IBM Spectrum Protect Plus.

Before you begin

Tip:

For information about compatible object storage providers, see [technote 108714](#).

Configure the key that is required for the cloud object. For instructions, see [“Adding an access key”](#) on page 121.

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 menu, click **System Configuration > Backup Storage > Object Storage**.
2. Click **Add Object Storage**.
3. From the **Provider** list, select **S3 Compatible Storage**.
4. Complete the fields in the **Object Storage Registration** pane:

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Enter the endpoint of 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. Click **Get Buckets**, and then select a bucket to serve as the target.

After the buckets are generated, the **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.

6. In the **Standard object storage bucket** field, select a bucket to serve as the backup target.
7. Optional: In the **Archive object storage bucket** field, 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 [“Copy snapshots to secondary backup storage” on page 6](#).

8. Click **Register**.

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” on page 145 . To modify an existing SLA policy, see “Editing an SLA policy” on page 149 .

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 menu, click **System Configuration > Backup Storage > Object Storage**.
2. Click the edit icon  that is associated with an object storage provider.
The **Update Object Storage** pane is displayed.
3. Revise the settings for the cloud provider, and then click **Update**.

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.

Procedure

To delete a cloud storage provider, complete the following steps:

1. In the navigation menu, click **System Configuration > Backup Storage > Object Storage**.
2. Click the delete icon  that is associated with a provider.
3. Click **Yes** to delete the provider.

Managing repository server storage

You can copy data to a repository server for longer-term data protection. For the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server Version 8.1.7 or later. To copy data to tape, IBM Spectrum Protect server Version 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 overview

To copy IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, you must complete configuration tasks in both environments.

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 process requests for backup and restore operations. The Amazon Simple Storage Service (S3) protocol enables communication between the two servers.

Action	How to
<p>Create a storage pool or pools for the data that is copied from IBM Spectrum Protect Plus.</p> <p>You can copy data to standard object storage or to tape storage.</p> <p>Copying data to standard object storage</p> <p>When data is copied to standard object storage, a full copy is created during the first copy operation. Subsequent copies are incremental and capture cumulative changes since the last copy operation. 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.</p> <p>To copy data to standard object storage, you must create a cloud-container or directory-container storage pool.</p> <p>Copying data to tape</p> <p>When data is copied to tape, a full copy of the data is created. Copying data to tape 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.</p> <p>To copy data to tape, you must create a cloud-container or directory-container storage pool and a cold-data-cache storage pool. A cloud-container or directory-container storage pool is required to store metadata that is used for restore and other IBM Spectrum Protect Plus operations.</p>	<p>To create storage pools by using the IBM Spectrum Protect Operations Center, following the instructions in the following topics:</p> <ul style="list-style-type: none"> • Directory-container: Configuring a directory-container storage pool for data storage • Cloud-container: Configuring a cloud-container storage pool for data storage • Cold-data-cache: Configuring operations for copying data to tape <p>Tip: Alternatively, issue the DEFINE STGPOOL command to create a storage pool as described in the following topics:</p> <ul style="list-style-type: none"> • Directory-container: Define a directory-container storage pool • Cloud-container: Define a cloud-container storage pool • Cold-data-cache: Define a primary storage pool for offloading data to tape
<p>Create a policy domain that points to the storage pool or pools.</p> <p>The policy domain defines the rules that control the backup services for IBM Spectrum Protect Plus.</p>	<p>To create a policy domain by using the Operations Center, follow the instructions in Creating a policy domain.</p> <p>Tip: Alternatively, issue the DEFINE OBJECTDOMAIN command to create a storage pool as described in Define a policy domain for object clients.</p>

Action	How to
<p>Add an object agent and object client on the IBM Spectrum Protect server.</p> <p>The object agent provides a gateway between the IBM Spectrum Protect Plus server and the IBM Spectrum Protect server.</p> <p>The object client identifies the IBM Spectrum Protect Plus server and enables IBM Spectrum Protect Plus to store objects in the IBM Spectrum Protect server.</p> <p>You must create the object agent before you can create the object client.</p>	<p>To create an object agent and object client, follow the instructions in Configuring an object agent service.</p>

Tasks for configuring IBM Spectrum Protect Plus

You must add the IBM Spectrum Protect server as a backup storage provider for IBM Spectrum Protect Plus. Then select the IBM Spectrum Protect Plus server as a data copy target in the service level agreement (SLA) policy for the resources that you want to protect.

Action	How to
Add the IBM Spectrum Protect server as a backup storage provider.	To add the IBM Spectrum Protect server to IBM Spectrum Protect Plus, follow the instructions in “Adding a repository server as a backup storage provider” on page 119.
Select the IBM Spectrum Protect server as a target for standard object storage or archive object storage (tape) in the SLA policy for the resources that you want to protect.	To create an SLA policy that defines the IBM Spectrum Protect server as backup storage target, follow the instructions in “Creating an SLA policy” on page 145.

Adding a repository server as a backup storage provider

Add 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 121 and [“Adding a certificate”](#) on page 122.

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.

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 an IBM Spectrum Protect server as backup storage provider complete the following steps:

1. In the navigation menu, click **System Configuration > Backup Storage > Repository Server**.
2. Click **Add Repository Server**.
3. Complete the fields in the **Register Repository Server** pane:

Name

Enter a meaningful name to help identify the repository server.

Hostname

Enter the high-level address (HLA) of the repository server object agent. Running the IBM Spectrum Protect `q serv OBJAGENT f=d` command retrieves this information.

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 help 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, then click **Upload**.

Copy and paste

Select to enter the name of the certificate, copy and paste the contents of the certificate, then click **Create**.

Use existing

Select to use a previously uploaded certificate.

4. Click **Register**.

The IBM Spectrum Protect server is added to the repository 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” on page 145. To modify an existing SLA policy, see “Editing an SLA policy” on page 149.

Related concepts

[“Configuration overview”](#) on page 117


To copy IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, you must complete configuration tasks in both environments.

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 navigation menu, click **System Configuration > Backup Storage > Repository Server**.
2. Click the edit icon  that is associated with a repository server provider.
The **Update Repository Server** pane is displayed.
3. Revise the settings for the repository server provider, and then click **Update**.

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.

Procedure

To delete a repository server provider, complete the following steps:

1. In the navigation menu, click **System Configuration > Backup Storage > Repository Server**.
2. Click the delete icon  that is associated with a repository server provider.
3. Click **Yes** to delete the provider.

Managing keys and certificates

Cloud resources and repository servers 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 user name and password of your copy destinations and allow them to be accessed by IBM Spectrum Protect Plus. Some copy destinations also require certificates for additional data security.

When utilizing a resource in IBM Spectrum Protect Plus that requires credentials to access a copy destination, select **Use existing key** or **Use existing certificate**, and select the associated key or certificate.

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

Add an SSH key to provide credentials for Linux-based resources, including file indexing and restore operations on virtual machines under vCenter and Hyper-V, as well as Oracle, Db2, and MongoDB application servers. SSH keys provide a secure connection between your resources and IBM Spectrum Protect Plus.

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.
- Ensure the public SSH key is placed in the appropriate `authorized_keys` file for the IBM Spectrum Protect Plus agent user. Typically, the file is located at `/home/<username>/.ssh/authorized_keys`. The `.ssh` directory and all files under it must have their permissions set to 600.

Procedure

To add a key, complete the following steps:

1. On your resource, generate an SSH key. For example, on an Oracle server, enter the `ssh-keygen` command and follow the instructions.
2. When prompted to Enter file in which to save the key, enter a file and location, for example: `/root/sshkey`.
3. In the `/root` location on the server entered in step 2, the file `sshkey.pub` contains the public key. This will later be copied, pasted, and saved into the `authorized_keys` file after executing `cd ~/.ssh` while logged in as the user assigned to IBM Spectrum Protect Plus.
4. In the IBM Spectrum Protect Plus navigation pane, click **System Configuration > Keys and Certificates**.
5. From the **SSH Keys** section, click **Add SSH Key**.
6. Complete the fields in the **SSH Key Properties** pane:

Name

Enter a meaningful name to help identify the SSH key.

User

Enter the user associated with the resource and SSH key.

Private key

Copy and paste the private key, which can be found in the `sshkey` file.

7. Click **Save**.


The key displays in the **SSH Keys** table and can be selected when utilizing a feature that requires credentials to access a resource through 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 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, a secondary site, and a demo site.

Adding a site

After you add a site to IBM Spectrum Protect Plus, you can assign backup storage servers to the site.

Procedure

To add a site, complete the following steps:

1. In the navigation pane, click **System Configuration > Site**.
2. Click **Add Site**.
The **Site Properties** pane is displayed.
3. Enter a site name.
4. Optional: To manage the network activity on a defined schedule, change the throughput for site replication and copy operations:
 - a) Select the **Enable Throttle** check box.
 - b) In the **Rate** field, adjust the throughput:
 - 1) Change the numerical rate of throughput by clicking the up or down arrows.
 - 2) Select a unit for the throughput. The choices include **bytes/s**, **KB/s**, **MB/s**, and **GB/s**.

The default throughput is 100 MB/s (megabytes per second).

Figure 15. Enabling different rates of throttling for different times to improve throughput

- c) In the weekly schedule table, select daily times for throttling, or select specific days and times for throttling.

Tip: To select a time, click a timeslot in the table. The selected timeslot is highlighted. To clear a timeslot, click a highlighted time slot. To select the same timeslot for every day of the week, click a timeslot in the **All** row.

After you make your selections, throttling days and times are listed underneath the schedule table.

5. Click **Save** to commit the changes and close the pane.

Results


The site is displayed in the sites table and can be applied to new and existing backup storage servers.

Editing a site

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 pane, click **System Configuration > Site**.
2. Click the edit icon  that is associated with a site.
The **Site Properties** pane is displayed.
3. Revise the site name.
4. Optional: To manage the network activity on a defined schedule, change the throughput for site replication and copy operations:
 - a) Select the **Enable Throttle** check box.
 - b) In the **Rate** field, adjust the throughput:
 - 1) Change the numerical rate of throughput by clicking the up or down arrows.

2) Select a unit for the throughput. The choices include **bytes/s**, **KB/s**, **MB/s**, and **GB/s**.
The default throughput is 100 MB/s (megabytes per second).

Site

Site Properties

Name:

☒ Enable Throttle

Rate:

Schedule

	12	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10	11	12
All																									
Sunday																									
Monday																									
Tuesday																									
Wednesday																									
Thursday																									
Friday																									
Saturday																									

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

Figure 16. Enabling different rates of throttling for different times to improve throughput

- c) In the weekly schedule table, select daily times for throttling, or select specific days and times for throttling.

Tip: To select a time, click a timeslot in the table. The selected timeslot is highlighted. To clear a timeslot, click a highlighted time slot. To select the same timeslot for every day of the week, click a timeslot in the **All** row.

After you make your selections, throttling days and times are listed underneath the schedule table.


5. 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 pane, click **System Configuration > Site**.
2. Click the delete icon  that is associated with a site.
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 371](#)

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 360](#)

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 allow 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 pane, click **System Configuration > LDAP/SMTP**.
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 SSL connections or 636 for SSL connections.

SSL

Enable the SSL 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 371](#)

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.

Procedure

To add an SMTP server, complete the following steps:

1. In the navigation pane, click **System Configuration > LDAP/SMTP**.
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-SSL connections or 443 for SSL 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.

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 360](#)


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**.
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**.
2. Click the delete icon  that is associated with the server.
3. Click **Yes** to delete server.

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.



Attention: Modify global preferences only if absolutely necessary. The modification of global preferences can affect your storage environment.

About this task

The **Global Preferences** pane lists configurable global parameter settings in the following categories: Application, General, Job, Logging, Protection, and Security. As the administrator, you can edit the settings to meet your organization's requirements. Any changes that you make to parameter default values apply to all IBM Spectrum Protect Plus operations when you save the changes.

Procedure

To edit the values for any setting and apply them globally, complete the following steps:

1. In the navigation pane, click **System Configuration > Global Preferences**.
2. To apply global application preferences, edit the settings in the Application category. The default values for the preferences are shown in the following image:

Global Preferences	
⚠ Do not change these preferences unless directed by IBM Support.	
Application	
Enable SQL Server databases restored in test mode eligible for backup	<input type="checkbox"/>
Maximum volume size for backup target LUNs on Windows (TB)	256
Maximum concurrent servers running backups	0
Perform DB backup when log backup chain is broken	<input type="checkbox"/>

You can edit the following application preferences:

Enable SQL Server databases restored in test mode eligible for backup

Back up SQL Server databases that were restored in test mode. When this option is selected, SQL Server databases that were restored in test mode are available for selection in the SQL Backup pane or ad hoc backup wizard.

Maximum volume size for backup target LUNs on Windows (TB)

The maximum size of the storage for a backup target.

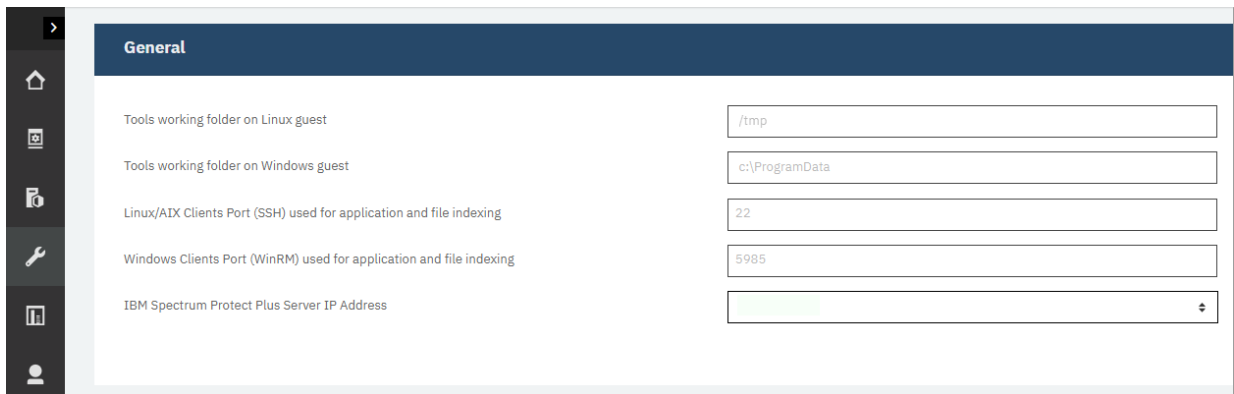
Maximum concurrent servers running backups

The maximum number of concurrent application servers per backup session.

Perform DB backup when 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.

3. To apply general preferences, edit the settings in the General category. The default values for the preferences are shown in the following image:



General	
Tools working folder on Linux guest	/tmp
Tools working folder on Windows guest	c:\ProgramData
Linux/AIX Clients Port (SSH) used for application and file indexing	22
Windows Clients Port (WinRM) used for application and file indexing	5985
IBM Spectrum Protect Plus Server IP Address	[Green Highlighted Address]

You can edit the following general preferences:

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.

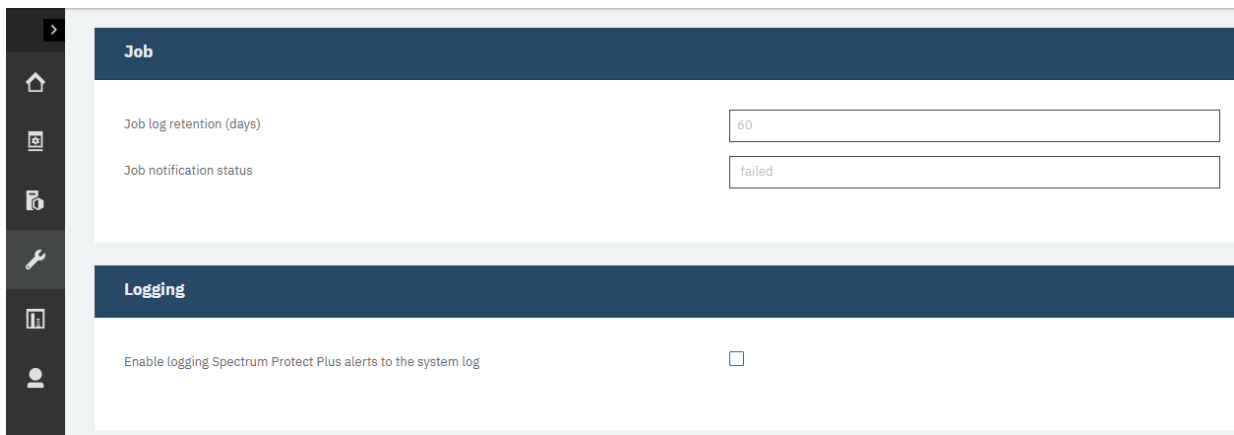
Windows Clients Port (WinRM) used for application and file indexing

The Windows Remote Management (WinRM) port that is used for application and file indexing on Windows clients.

IBM Spectrum Protect Plus Server IP Address

The list of available IP addresses for the IBM Spectrum Protect Plus server.

- To apply job or logging preferences, edit the values in the Job or Logging categories. The default values for the preferences are shown in the following image:



Job	
Job log retention (days)	60
Job notification status	failed

Logging	
Enable logging Spectrum Protect Plus alerts to the system log	<input type="checkbox"/>

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.

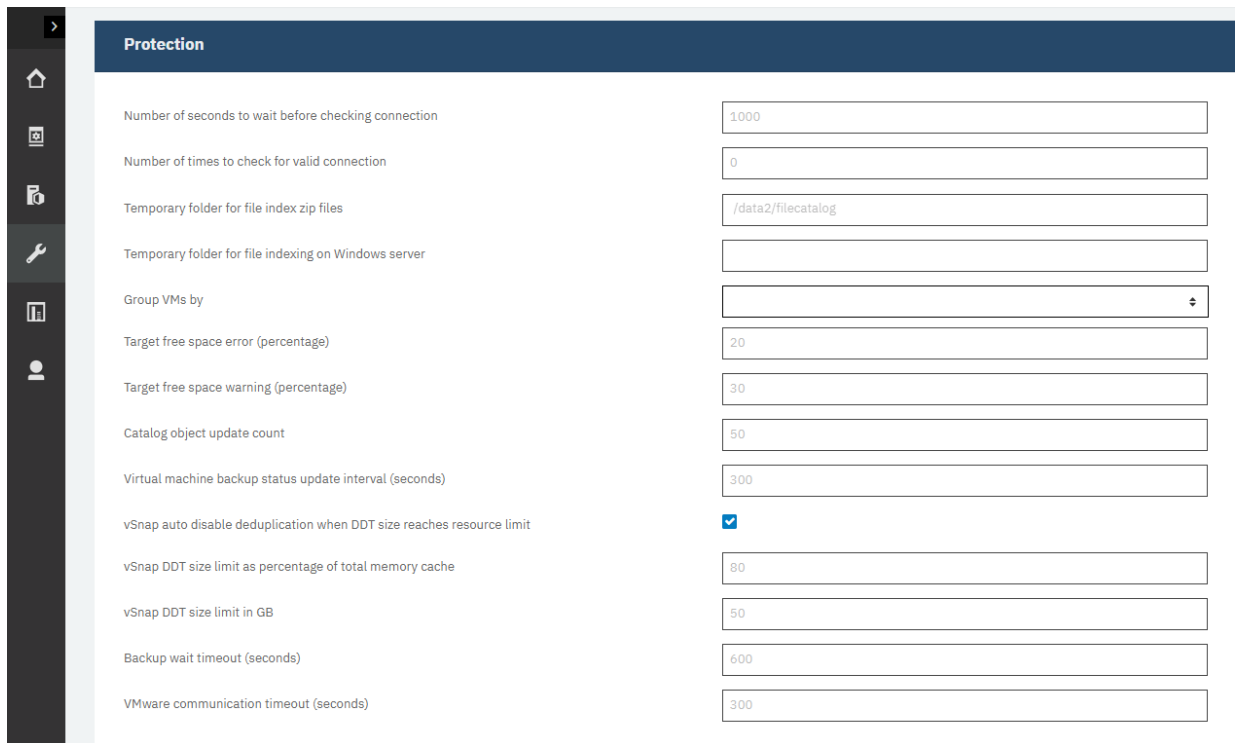
Job notification status

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.

- To apply protection preferences, edit the settings in the Protection category. The default values for the preferences are shown in the following image:



Setting	Value
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	
Target free space error (percentage)	20
Target free space warning (percentage)	30
Catalog object update count	50
Virtual machine backup status update interval (seconds)	300
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
Backup wait timeout (seconds)	600
VMware communication timeout (seconds)	300

You can edit the following protection preferences:

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. The group can be defined by a count of the VMs that are included in the group or the size of the VMs that are included in the 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.

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.

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.

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.

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.

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 connection 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.

6. To apply a security preference, edit the setting in the Security category. The default value for the preference is shown in the following image:

A screenshot of the IBM Spectrum Protect Plus configuration interface. It shows a 'Security' category header. Below it, there is a setting 'Set Minimum Password Length (characters)' with a text input field containing the value '8'.

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.

Customizing global preferences

As the administrator, you can modify the global preferences that are displayed in the Global Preferences pane as described in the topic “Configuring global preferences” on page 130. If a preference is not available in the Global Preferences pane, you can customize additional preferences manually. After you add a global preference by modifying the appropriate configuration files, you can update the preference in the Global Preferences pane.

Before you begin

You must have administrative access to the IBM Spectrum Protect Plus virtual appliance to add global preferences.



Attention: Modification of the global preferences properties file and associated JSON files can affect your storage environment. Modify global preferences only if absolutely necessary.

About this task

By default, the Global Preferences pane lists configurable global parameter settings in the following categories: Application, General, Job, Protection, and Security. The Global Preferences pane can be modified to display additional global parameter settings.

The preference schema and preference names are defined in the `preferences_en_us.properties` file, which is located in the `/etc/config/messages` directory. An administrator can add or modify the preference schema information and names in this file. The preference name property must be defined in this file. The preference name key must be unique.

The `preferences_en_us.properties` file is similar to the following example:

```
pref.category.api=api

pref.category.application=application
pref.category.catalog=catalog
pref.category.general=general
pref.category.job=job
pref.category.protection=protection
pref.category.recovery=recovery
pref.category.storage=storage

pref.subtype.application=application
pref.subtype.communication=communication
pref.subtype.common=common
pref.subtype.condense=condense
pref.subtype.configuration=configuration
pref.subtype.connection=connection
pref.subtype.file=file
pref.subtype.history=history
pref.subtype.hypervisor=hypervisor
pref.subtype.log=log
pref.subtype.maintenance=maintenance
pref.subtype.misc=misc
pref.subtype.monitor=monitor
pref.subtype.notification=notification
pref.subtype.replication=replication
pref.subtype.resource=resource
pref.subtype.schedule=schedule
pref.subtype.stat=statistics
pref.subtype.vmware=vmware

pref.security.user.passwordrule=Password Rules
pref.security.user.passwordrule.minimumLength=Minimum Length
pref.security.user.passwordrule.minimumLength_DETAIL=Minimum password length

pref.job.schedule.maintenanceSchedule=Maintenance job schedule
pref.job.notification.status=Job notification status
pref.job.log.retention.days=Job log retention (days)
pref.protection.vsphere.communication.timeout.seconds=VMware communication timeout (seconds)
pref.protection.protectionprovider.targetfreespacewarnpercentage=Target free space warning (percentage)
pref.protection.protectionprovider.targetfreespaceerrorpercentage=Target free space error (percentage)
pref.protection.protectionprovider.updateinterval=Update interval
pref.protection.protectionprovider.waitforothersla.enablebackupwait=Two slas backing up same vm
pref.protection.protectionprovider.waitforothersla.backupwaittimeout=Backup wait timeout
pref.application.common.iscsi.maxvolumesize=Maximum volume size
pref.application.common.iscsi.minvolumesize=Minimum volume size
pref.application.protection.blbackup.maxconcurrentservers=Maximum concurrent servers running backups
```

The preference information is defined in a group of preference files in JSON format located on the IBM Spectrum Protect Plus virtual appliance in the `/etc/config/preferences` directory. To define a preference, you must specify at least the preference name, type, and `defaultValue`. If other preference information is not provided, default values are applied. The `values` array holds information that is contained for drop-down selection and related subfields. Each object in the `values` array has information related to the subfield that is associated with the drop-down option. Each object inside the `values` array also has two additional properties called `optionName` and `option`.

In the sample `security.json` file, all of the required values are defined:

```
[
```

```
{
  "name": "pref.security.user.passwordrule",
  "category": "pref.category.security",
  "subtype": "pref.subtype.user",
  "type": "string",
  "isArray": true,
  "defaultValue": ["Minimum Length"],
  "uiOption": true
},
{
  "name": "pref.security.user.passwordrule.minimumLength",
  "category": "pref.category.security",
  "subtype": "pref.subtype.user",
  "type": "integer",
  "defaultValue": 8
}
]
```

Procedure

1. Connect to the IBM Spectrum Protect Plus virtual appliance by using the Secure Shell (SSH) network protocol.
2. Navigate to the following directory:
/etc/config/messages
3. Edit the preferences_en_us.properties file by using the vi text editor and modify the configuration parameters by providing the appropriate values for your environment.
4. Save the preferences_en_us.properties file.
5. Navigate to the directory that contains the JSON files that are associated with global preferences:
/etc/config/preferences
6. Edit the appropriate JSON file by using the vi text editor and modify the configuration parameters by providing the appropriate values for your environment.
7. Save the JSON file.
8. To view the newly defined global preference, go to the navigation pane and click **System Configuration > Global Preferences**.

What to do next

Optionally, you can update the newly defined global preference by completing the following steps:

1. In the navigation pane, click **System Configuration > Global Preferences**.
2. Modify the defined value for the preference and click outside of the modified value.

A checkmark and the word **Updated** indicate that the preference is updated.

Related tasks

[“Configuring global preferences” on page 130](#)

As the administrator, you can configure preferences that apply to all IBM Spectrum Protect Plus operations in the **Global Preferences** pane.

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:

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, select one of the following authentication types in the **Authentication Type** list:

Authentication Type	Logon information
IBM Spectrum Protect Plus	To log on as an IBM Spectrum Protect Plus user with SYSADMIN privileges, enter your administrator user name and password.
System	To log on as a system user, enter the serveradmin password. 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 91 .

What to do next

Review the configuration of the IBM Spectrum Protect Plus virtual appliance.

Related concepts

[“System requirements” on page 11](#)

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 367](#)

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.

Setting the time zone

Use the Administrative Console to set the time zone of the IBM Spectrum Protect Plus 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, select one of the following authentication types in the **Authentication Type** list:

Authentication Type	Login information
IBM Spectrum Protect Plus	To log in as an IBM Spectrum Protect Plus user with SYSADMIN privileges, enter your administrator user name and password.
System	To login as a system user, enter the serveradmin password. 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 91 .

3. Click **Perform System Actions**.
4. In the **Change Time Zone** 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 appliance when logged in with the user ID **serveradmin**.
5. Restart the IBM Spectrum Protect Plus appliance from the Administrative Console.
6. Once the IBM Spectrum Protect Plus 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.

Uploading an SSL certificate from the administrative console

To establish secure connections in IBM Spectrum Protect Plus, you can upload an SSL certificate such as an HTTPS or LDAP certificate by using the administrative console.

About this task

For HTTPS certificates, PEM encoded certificates with `.cer` or `.crt` extensions are supported.

For LDAP/Hyper-V certificates, DER encoded certificates with `.cer` or `.crt` extensions are supported. If you are uploading an LDAP SSL certificate, ensure that IBM Spectrum Protect Plus has connectivity to the LDAP server and that the LDAP server is running.

ASCII and binary format certificates are accepted with the standard `.pem`, `.cer`, and `.crt` file extensions. However, the administrative console certificate import function cannot be used to update the appliance SSL web server communications. To upload ASCII and binary format certificates, use the command line as described in [“Uploading an SSL certificate from the command line” on page 138](#)

Procedure

To upload an SSL certificate, complete the following steps:

1. Contact your network administrator for the name of the certificate to export.
2. From a supported browser, export the certificate to your computer. Make note of the location of the certificate on your computer. The process of exporting certificates varies based on your browser.
3. 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.

4. In the logon window, select one of the following authentication types in the **Authentication Type** list:

Authentication Type	Logon information
IBM Spectrum Protect Plus	To log on as an IBM Spectrum Protect Plus user with SYSADMIN privileges, enter your administrator user name and password.
System	To log on as a system user, enter the serveradmin password. 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 91 .

5. Click **Manage your certificates**.

6. Click **Browse**, and select the certificate that you want to upload.
7. Click **Upload SSL certificate for HTTPS**.
8. Restart the virtual machine where the application is deployed.

Uploading an SSL certificate from the command line

To upload ASCII and binary format certificates, use the command line for the IBM Spectrum Protect Plus virtual appliance. Certificates are accepted with the standard .pem, .cer, and .crt file extensions.

About this task

This process requires that you package the private key, public key, and chain certificates into a PKCS12 format file (often referred to as PFX file with .p12 extension) and import this manually into the IBM Spectrum Protect Plus Java keystore. The procedure assumes you already have the private, public, and all supporting security objects provided by your security vendor packaged into a PKCS12 format file named *name*.p12.

If you do not have this file, you must work with your security vendor using a separate server and/or OpenSSL to generate the necessary certificate signing request. Once received, package the resulting private, public, and chain certificate objects into the required *name*.p12 file.

Procedure

To import the *name*.p12 file, complete the following steps:

1. Log on with the user ID **serveradmin** on the IBM Spectrum Protect Plus virtual appliance.
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 91](#).
2. At the command line execute the following command:

```
/usr/java/latest/bin/keytool -importkeystore -deststorepass ecx-beta -  
destkeystore /opt/virgo/configuration/keystore -srckeystore NAME.p12 -  
srcstoretype PKCS12
```
3. Restart the virtual appliance.

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 91](#).

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 91](#).

What to do next

Enter commands to administer the virtual appliance. To log off, type `exit`.

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: https://www.ibm.com/support/knowledgecenter/en/STXKQY_5.0.3/com.ibm.spectrum.scale.v5r03.doc/bl1adv_firewallportopenexamples.htm.

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
```

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.

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 **pvdisk** 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 91](#).

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.

```
echo "- - -" > /sys/class/scsi_host/host#/scan
```

Where # is the latest host number.

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 91](#).

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:

```
[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 **pvdisk** command.
4. To create a new Physical Volume (PV), issue the **pvcreeate /dev/sdd1** command.
5. To view the new PV from /dev/sdd1, issue the **pvdisk** command.
6. To review the Volume Group (VG), issue the **vgdisk** 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 **vgdisk** 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
```

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 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.

Silver

This policy runs daily with a retention period of 1 month.

Bronze

This policy runs daily with a retention period of 1 week.

To view and manage backup policies and to monitor the virtual machines and databases that are protected by policies, click **Manage Protection > Policy Overview** in the navigation pane.

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.

For installations of IBM Spectrum Protect Plus V10.1.5, a demo SLA configuration is available for testing. This demonstration feature includes the following elements:

- A demonstration site named **Demo**
- An SLA policy named **Demo**
- A local vSnap configuration for the demo SLA.

You can choose to use the demo site for testing backup and restore operations. The data is backed up to the local vSnap configuration when you run the demo SLA policy.

Note: The built-in vSnap is set so that it can be used only by the Demo Site. Do not use the built-in IBM Spectrum Protect Plus vSnap with any other site.

Creating an SLA policy

You can create custom SLA policies to define backup frequency, retention, replication, and copy policies that are specific for your environment.

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, complete the following steps:

1. In the navigation pane, click **Manage Protection > Policy Overview**.
2. Click **Add SLA Policy**.
The **New SLA Policy** pane is displayed.
3. In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
4. In the **Operational Protection** section under **Main Policy**, set the following options for backup operations. These operations occur on the vSnap servers that are defined in the **System Configuration > Backup Storage > Disk** window.

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

Select this check box to create the main policy without defining a frequency or start time. Policies created without a schedule can be run on-demand.

Frequency

Enter a frequency for backup operations.

Start Time

Enter the date and time that you want the backup operation to start.

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 check box 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.

5. 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 [“Establishing a replication partnership for a vSnap server” on page 82](#).

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this check box to create the replication relationship without defining a frequency or start time.

Frequency

Enter a frequency for replication operations.

Start Time

Enter the date and time that you want the replication operation to start.

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.

6. 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 check box to create the copy relationship without defining a frequency or start time.

Frequency

Enter a frequency for copy operations.

Start Time

Enter the date and time that you want the copy operation to start.

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: Copy retention options are disabled if a server that uses write once read many (WORM) retention is selected in the **Target** field.

Source

Click the source for the copy operation:

Main Policy Destination

The source for the copy operation is the target site that is defined in the **Main 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 secondary backup storage” on page 111](#) 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 check box to create the archive relationship without defining a frequency or start time.

Frequency

Enter a frequency for archive operations.

Start Time

Enter the date and time that you want the archive operation to start.

Retention

Specify the retention period for the archive snapshots as a unit of time in days, months, or years.

Source

Click the source for the archive destination:

Main Policy Destination

The source for the archive operation is the target site that is defined in the **Main 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 111](#).

7. Click **Save**. The SLA policy can now be applied to backup job definitions.

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 368
Create a backup job definition that uses the SLA policy.	See the backup topics in Chapter 9, “Protecting hypervisors,” on page 151 and Chapter 10, “Protecting applications,” on page 189 .

Related concepts

[“Replicate backup-storage data ” on page 5](#)

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.

[“Copy snapshots to secondary backup storage” on page 6](#)


The vSnap server is the primary backup location for snapshots. All IBM Spectrum Protect Plus environments have at least one vSnap server. Optionally, you can copy snapshots from a vSnap server to secondary backup storage.

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 pane, click **Manage Protection > Policy Overview**.
2. Click the edit icon  that is associated with a policy.
The **Edit SLA Policy** 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.

Procedure

To delete an SLA policy, complete the following steps:

1. In the navigation pane, click **Manage Protection > Policy Overview**.
2. Click the delete icon  that is associated with an SLA policy.
3. Click **Yes** to delete the policy.
4. If you are deleting the demo SLA policy, go to **System Configuration > Site**, and delete the site named Demo.

Note:

When you delete the demo site you must register the local host vSnap with user credentials to another valid site.

Chapter 9. Protecting hypervisors

You must register the hypervisors that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the virtual machines and resources that are associated with the hypervisors.

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.

System requirements

Ensure that your VMware environment meets the system requirements in [“Hypervisor requirements”](#) on page 26.

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](#).

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 pane, click **Manage Protection > Hypervisors > 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. Select the **Use SSL** check box to enable an encrypted Secure Sockets Layer (SSL) connection. The typical default port is 80 for non SSL connections or 443 for SSL connections.

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.

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 368.

Related concepts

[“Managing identities”](#) on page 373

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 155

Use a backup job to back up VMware resources such as virtual machines, datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data”](#) on page 164

VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Virtual machine 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.

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, refer to the VMware documentation about adding a permission to an inventory object.

vCenter Server Object	Required Privileges
Alarm	<ul style="list-style-type: none"> Acknowledge alarm Set alarm status

vCenter Server Object	Required Privileges
Cryptographic Operations (6.5 and 6.7)	<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
Host > Configuration	<ul style="list-style-type: none"> • Storage partition configuration
Inventory Service > Tagging (6.0) vSphere Tagging (6.5 and 6.7)	<ul style="list-style-type: none"> • Assign or Unassign vSphere Tag • 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

vCenter Server Object	Required Privileges
Virtual Machine > Configuration	<ul style="list-style-type: none"> • Add existing disk • Add new disk • Add or remove device • Advanced (6.0 and 6.5) • Advanced configuration (6.7) • Change CPU count • Change memory (6.7) • Configure raw device (6.7) • Disk change tracking (6.0 and 6.5) • Memory (6.0 and 6.5) • Modify device settings • Raw device (6.0 and 6.5) • Reload from path • Remove disk • Rename • Settings • Toggle disk change tracking (6.7)
Virtual Machine > Guest Operations	<ul style="list-style-type: none"> • Guest Operation Modifications • Guest Operation Program Execution • Guest Operation Queries
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"> • Mark as template • Mark as virtual machine
Virtual Machine > Snapshot management	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert snapshot

vCenter Server Object	Required Privileges
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

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.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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 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 pane, click **Manage Protection > Hypervisors > 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 **Use existing user**.
5. Select a user in the **Select user** list.
6. Click **Test**.

Backing up VMware data

Use a backup job to back up VMware resources such as virtual machines, 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 151.

- Configure SLA policies. For more instructions, see [“Create backup policies” on page 93](#).
- 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 15, “Managing user access,” on page 363](#).
- If a virtual machine 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 virtual machine, 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 virtual machines.

About this task

- When backing up VMware virtual machines, 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.

Procedure

To define a VMware backup job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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 virtual machines.

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 349](#).

When 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. 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 159](#).

5. To edit options before you create the job definition, 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 virtual machine snapshot. All VSS-compliant applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL, and the system state are quiesced. VMDKs and virtual machines 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 virtual machine 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 during the backup job, enable the **Truncate SQL logs** option. The credentials must be established for the associated virtual machine by using the Guest OS user name and Guest OS Password option within the backup job definition. When the virtual machine 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 virtual machine by using an SSH key, or the **Guest OS Username** and **Guest OS Password** options within the backup job definition. Ensure that the virtual machine 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.

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.

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 virtual machine. Enter the user name and password, and ensure that the virtual machine 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 virtual machine, use the **Test** function.

The **Test** function verifies communication with the virtual machine and tests DNS settings between the IBM Spectrum Protect Plus appliance and the virtual machine. To test a connection, select a single virtual machine, and then click **Select Options**. Select **Use existing user** and select a previously entered user name and password for the resource, and then click **Test**.

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. 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.

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.

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 161 .
Create a VMware restore job definition.	See “Restoring VMware data” on page 164 .

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 350](#)

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 345](#)

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 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:

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware**.

2. Select **VMs and Templates** in the **View** filter.
3. Click the link for the vCenter, and then click the link for the virtual machine that contains the VMDKs that you want to exclude.
4. Select one or more VMDKs, and then click **Select SLA Policy**.
5. Clear the check box for the selected SLA policy, and then click **Save**.

Backing up a Linux-based vCenter Server Appliance

To back up a Linux-based vCenter Server Appliance, you must modify the VMware pre-freeze and post-thaw scripts on the vCenter virtual machine to avoid corrupted vCenter backups.

Procedure

To modify the scripts, complete the following steps:

1. On the virtual machine, navigate to the `/usr/sbin` directory and replace the content of the `pre-freeze-script` script with the following content:

```
#!/bin/bash
#set log directory
log="/var/log/vpostgres_backup.log"
#set and log start date
today=`date +%Y/%m/%d\ %H:%M:%S`
echo "${today}: Start of creation consistent state" >> ${log}
#execute freeze command
cmd="echo \"SELECT pg_start_backup('${today}', true);\" | sudo /opt/vmware/vpostgres/9.4/bin/psql -U postgres >> ${log} 2>&1"
eval ${cmd}
#set and log end date
today=`date +%Y/%m/%d\ %H:%M:%S`
echo "${today}: Finished freeze script" >> ${log}
```

2. Replace the content of the `post-thaw-script` script with the following content:

```
#!/bin/bash
#set log directory
log="/var/log/vpostgres_backup.log"
#set and log start date
today=`date +%Y/%m/%d\ %H:%M:%S`
echo "${today}: Release of backup" >> ${log}
#execute release command
cmd="echo \"SELECT pg_stop_backup();\" | sudo /opt/vmware/vpostgres/9.4/bin/psql -U postgres >> ${log} 2>&1"
eval ${cmd}
#set and log end date
today=`date +%Y/%m/%d\ %H:%M:%S`
echo "${today}: Finished thaw script" >> ${log}
```

Managing VADP backup proxies

In IBM Spectrum Protect Plus, you can 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)

If proxies exist, the entire processing load is shifted off the host system and onto the proxies. If proxies do not exist, the entire load stays on the host. Throttling ensures that multiple VADP proxies are optimally utilized to maximize data throughput. For each 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 no other proxies exist, the host takes over the job. If a proxy server becomes unavailable when a job is running, the job might 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.
- 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 [“VADP proxy requirements” on page 20](#).

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 369](#).

Tip: The IBM Spectrum Protect Plus version of the VADP proxy installer includes Virtual Disk Development Kit (VDDK) version 6.5. This version of the VADP proxy installer provides the external VADP proxy support with vSphere 6.5.

Procedure

To create VMware VADP proxies, complete the following steps:

1. In the navigation pane, 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.

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**.

The proxy is added to the **VADP Proxy** table.

5. Click **Register** to register the proxy server.

You can unregister or suspend the server by using 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, select **Actions > Resume**.

After successful registration, the service vadm is started on the proxy machine. A log file, vadm.log, is generated in /opt/IBM/SPP/logs directory.

6. Repeat the previous steps for each proxy you want to create.

The connection between the IBM Spectrum Protect Plus virtual appliance 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 the IBM Spectrum Protect Plus virtual appliance 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 vadm_ip`, where vadm_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 155 . The proxies are indicated in the job log by a log message similar to the following text: Run remote vmdbkbackup of MicroService: http://<proxy> nodename, IP:proxy_IP_address

Related tasks

[“Setting options for VADP proxies” on page 163](#)

When you create VADP proxies in IBM Spectrum Protect Plus, you can configure various options for each VADP proxy.

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 369](#).

Procedure

To set options for VMware VADP proxies, complete the following steps:

1. In the navigation pane, 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 **Set Options**.
4. Complete the following fields in the **Set VADP Proxy Options** pane:

Site

Assign a site to the proxy.

User

Select a previously entered user name for the provider.

Transport Modes

Set the transport modes to be used by the proxy. 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.

To turn off compression, select **disabled**.

Log retention in days

Set the number of days to retain logs before they are deleted.

Read and write buffer size

Set the buffer size of the data transfer, measured in bytes.

Block size of NFS volume

Set the block size to be used by the mounted NFS volume, measured in bytes.

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 CPU = 1 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.

What to do next

After you create the VADP proxies, you can complete the following actions:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 155.
Uninstall the proxies when you cease running the VMware backup jobs.	See “Uninstalling VADP proxies” on page 164.

Related tasks

[“Creating VADP proxies” on page 161](#)

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:

1. From a command prompt, navigate to the directory `/opt/IBM/SPP/uninstall` on the proxy host system.
2. Run the following command:
`./uninstall_vmdkbackup`

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 155.](#)
- 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 15, “Managing user access,” on page 363](#) and [“Managing user accounts” on page 371.](#)
- 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.
- If you plan to restore 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 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:

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. For instructions for creating a fenced network, see [“Creating a fenced network through a VMware restore job”](#) on page 170.

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.

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.

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 pane, click **Manage Protection > Hypervisors > VMware > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:


- You can also open the "Snapshot restore" wizard by clicking **Jobs and Operations > Create job > Snapshot restore > VMware**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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	<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: <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 > Site pane.</p>

Option	Description
	<p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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.

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 172 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 system 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 test, production, or clone mode. After the job is created, it can be run in production or clone mode through the **Job Sessions** pane. You can also change the name of the restored VM by entering the new VM name in the **Rename VM (optional)** field. Click **Next** to continue.
9. 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.

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

Enable this option to automatically clean up allocated resources as part of a restore job if the virtual machine recovery fails.

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.

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.

10. 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.

11. 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**.

12. 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 151](#)

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 164](#).

Procedure

To create a fenced network, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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 151

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 164](#) 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 172.
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 164.
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 173 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.

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 requirements”](#) on page 26.

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 pane, click **Manage Protection > Hypervisors > 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 SSL** check box to enable an encrypted Secure Sockets Layer (SSL) connection.

To enable an SSL connection, you must add the self-signed SSL certificate for the Hyper-V server or a certificate authority (CA) certificate. To upload a certificate, see [“Uploading an SSL certificate from the administrative console”](#) on page 137.

If you do not select **Use SSL**, you must complete additional steps on the Hyper-V server. See [“Enabling WinRM for connection to Hyper-V servers”](#) on page 175.

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 368 .

Related tasks

[“Backing up Hyper-V data” on page 176](#)

Use a backup job to back up Hyper-V data with snapshots.

[“Restoring Hyper-V data” on page 179](#)

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 SSL 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.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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 Hyper-V Server virtual machine. The test function 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 pane, click **Manage Protection > Hypervisors > 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 **Use existing user**.
5. Select a user in the **Select user** list.
6. 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 174](#)
- Configure SLA policies. For instructions, see [“Create backup policies” on page 93](#).
- 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](#).

- 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 15, “Managing user access,” on page 363](#).
- If a virtual machine 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.

Procedure

To define a Hyper-V backup job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > 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 349.

5. To edit options before you start the job, click the edit icon in the table **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 virtual machine snapshot.

VM Snapshot retry attempts

Set the number of times IBM Spectrum Protect Plus should attempt to snapshot a virtual machine 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. Note that credentials must be established for the associated virtual machine 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 virtual machine 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 virtual machine by using an SSH key, or a Guest OS Username and Guest OS Password option in the backup job definition. Ensure that the virtual machine 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.

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.

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 virtual machine. Enter the username and password, and ensure that the virtual machine 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 virtual machine, use the **Test** function.

The **Test** function verifies communication with the virtual machine and tests DNS settings between the IBM Spectrum Protect Plus appliance and the virtual machine. To test a connection, select a single virtual machine, then click **Select Options**. Select **Use existing user** and select a previously entered user name and password for the resource, and then click **Test**.

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 virtual machines 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 179 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 350](#)

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 345](#)

You can run any job on demand, even if the job is set to run on a schedule.

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 176](#).
- 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 15, “Managing user access,” on page 363](#) and [“Managing user accounts” on page 371](#).

- 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 using clone mode and using original IP configuration, ensure that credentials are established through the Guest OS Username and Guest OS Password option 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:

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.

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.

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 pane, click **Manage Protection > Hypervisors > Hyper-V > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Hyper-V**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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: <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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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, 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.

6. 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.

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 Hyper-V host or cluster, specify the following network settings:

Allow system 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 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**, select the restore method to be used for source selections. Set the Hyper-V restore job to run in test, production, or clone mode by default. After the job is created, you can run the job in production or clone mode by using the **Job Sessions** pane. You can also change the name of the restored VM by entering the new VM name in the **Rename VM (optional)** field. Click **Next** to continue.
9. 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 vMotion operation for the resources in the background. The destination of the vMotion 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.

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

Enable this option to automatically clean up allocated resources as part of a restore job if the virtual machine recovery fails.

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.

10. 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.

11. 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**.

12. 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 176](#)

Use a backup job to back up Hyper-V data with snapshots.

[“Adding a Hyper-V server” on page 174](#)

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.

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 27](#).
- 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 `xfsprogs` package is between 3.2.0 and 4.1.9, the file restore can fail due to a known issue in `xfsprogs` that causes corruption of a clone or snapshot file system when its UUID is modified. To resolve this issue, update `xfsprogs` 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 pane, click **Manage Protection > 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 391](#).
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.
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 155](#)

Use a backup job to back up VMware resources such as virtual machines, datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data” on page 164](#)

VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Chapter 10. Protecting applications

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.

Ensure that your Db2 environment meets the system requirements. For more information, see [“Db2 requirements”](#) on page 36.

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 192.

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.

Requirements for the IBM Spectrum Protect Plus Db2 application server are available here, [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 dedicated local database directories are allowed for each database in your volume setup.

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 192.
- 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.

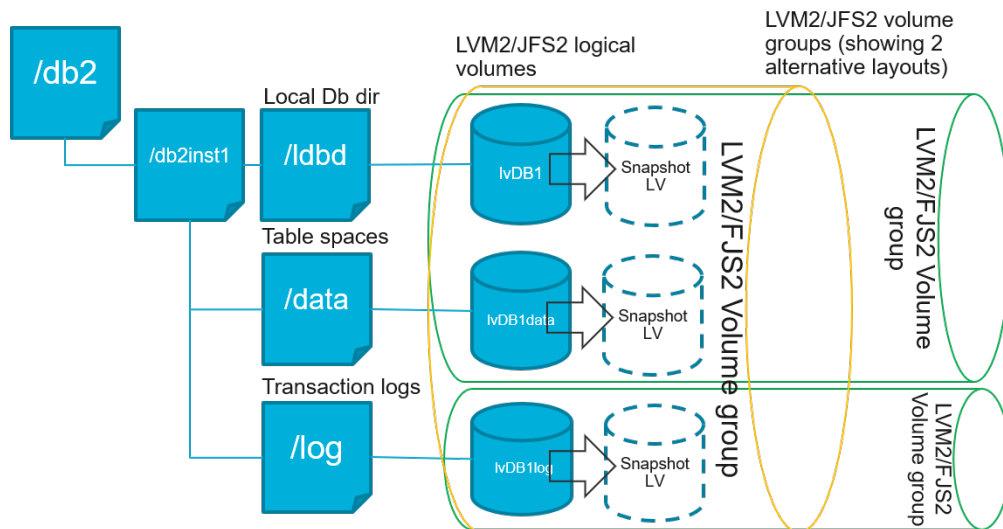


Figure 17. Logical volume layout examples




Figure 18. 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 `lv` 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 > Applications > Db2**.
2. In the **Db2** window, click **Manage application servers**, and click **Add application server** to add the host machine.



Figure 19. Adding a Db2 agent

3. In the **Application Properties** section, enter the host address.
4. 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.

A screenshot of the IBM Spectrum Protect Plus web interface. The left sidebar shows navigation icons. The main content area is titled "Db2" and contains a "Manage application servers" section. Within this section, there is an "Application Properties" form. The form includes a "Host Address" field with the value "77.00.999.12". Below this are two radio buttons: "User" (selected) and "SSH Key". Under the "User" option, there are three fields: "Use existing user" (a checkbox), "User ID" (containing "domain\user"), and "Password" (containing "Password"). At the bottom of the form are "Cancel" and "Save" buttons. A "Create job" button is visible in the top right corner of the interface.

Figure 20. 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.

5. 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](#).

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 pane, expand **Manage Protection > Applications > Db2**.

Tip: To add more Db2 instances to the **Instances** pane, follow the instructions in [Adding a Db2 application server](#).

2. Click **Run Inventory**.

The screenshot shows the 'Db2 Backup' section of the IBM Spectrum Protect Plus interface. It features a search bar, a dropdown for 'All Databases', and a table of Db2 instances. The first instance is selected, and the 'Run Inventory' button is highlighted.

Instances	Last Inventory at Oct 14, 2019 10:17:23 AM	Run Inventory	
Name	Hosts (partition numbers)	Version	SLA Policy
<input checked="" type="checkbox"/> oblong2.storage.ork.ibm.com/db2cat1	oblong2.storage.ork.ibm.com(0,1,2,3)	11.1.0	Gold
<input type="checkbox"/> oblong2.storage.ork.ibm.com/db2cat2	oblong2.storage.ork.ibm.com(0)	10.5.0	
<input type="checkbox"/> oblong2.storage.ork.ibm.com/db2cat3	oblong2.storage.ork.ibm.com(2,3)	11.1.0	
<input type="checkbox"/> oblong2.storage.ork.ibm.com/db2cat4	oblong2.storage.ork.ibm.com(0)	10.5.0	
<input type="checkbox"/> oblong2.storage.ork.ibm.com/db2cat5	oblong2.storage.ork.ibm.com(0)	11.1.0	

Total: 10

Buttons: Clear Selections (1), Run, Select an SLA policy, Select Options

Figure 21. Detecting Db2 resources

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 pane, click **Manage Protection > Applications > 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.

Test result of kasab5			
1. Physical - Basic Test for physical host network configuration			
Name	Description	Status	Message
Host FQDN Resolvable Test	Host FQDN must be resolvable to an IPv4 address	✓	
Socket Connection Test	Must allow socket connection on port 22 for Linux	✓	
2. Remote - Remote executor test for session creation and remote agent deployment			
Name	Description	Status	Message
Remote Session Test	Latest remote agent must be installed on host, SSH and SFTP service must be installed on Linux host, and port must be open to create session to SSH service.	✓	
Remote Agent Execute Test	Remote agent must be configured correctly using user credentials with sufficient privileges.	✓	
3. AIX - Basic AIX prerequisites for file and volume operations			
Name	Description	Status	Message
Sudo Privileges	User must have password-less sudo privileges	✓	

Figure 22. Testing the connection

The test report shows a list of the tests. It consists of a test for the physical host network configuration, and tests for the remote server installation on the host, which checks SSH and SFTP on the host. The third test checks for operating system prerequisites and correct sudo privileges.

- 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:

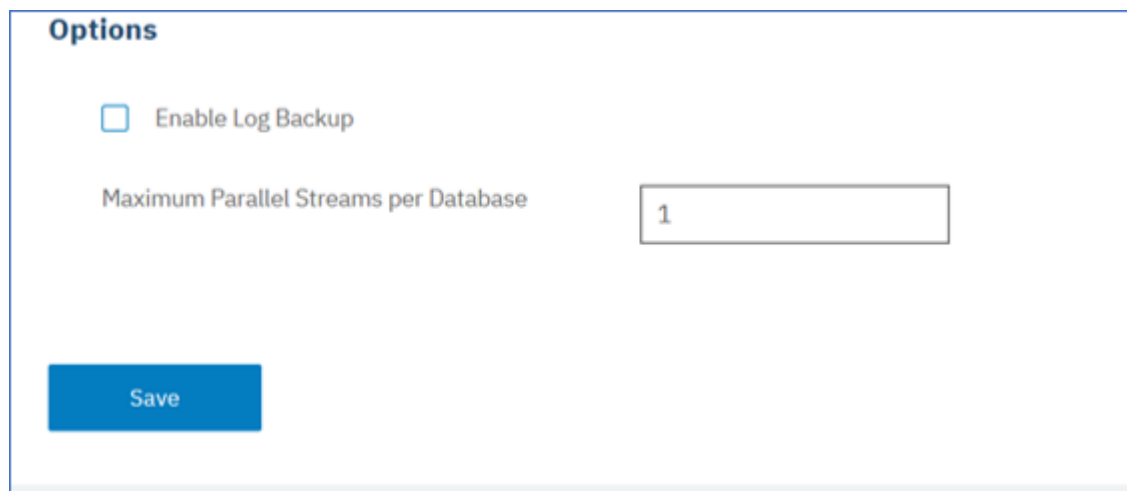
- Add the application servers that you want to back up. For the procedure, see [Adding a Db2 application server](#).
- Configure a Service Level Agreement (SLA) Policy. For the procedure, see [Defining a Service Level Agreement backup job](#).

- 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 [Managing user access](#).
- 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.

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > 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.

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](#).



The screenshot shows a user interface for configuring backup options. The title 'Options' is at the top left. Below it is a checkbox labeled 'Enable Log Backup'. Further down is a label 'Maximum Parallel Streams per Database' next to a text input box containing the value '1'. A blue 'Save' button is located at the bottom left of the form area.

Figure 23. Backup pane with the Enable Log Backup option

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.

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 > Applications > 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.

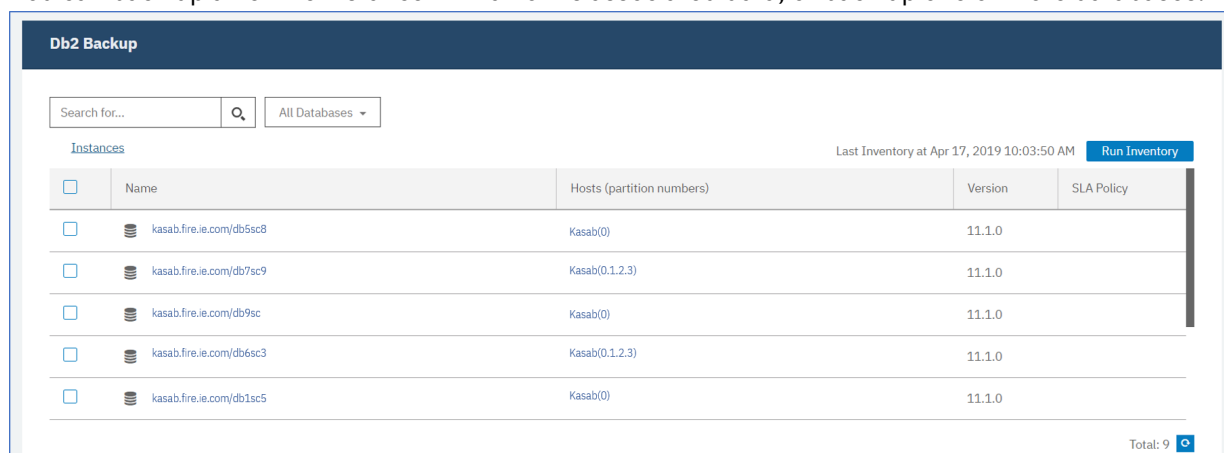


Figure 24. Db2 Backup pane showing instances

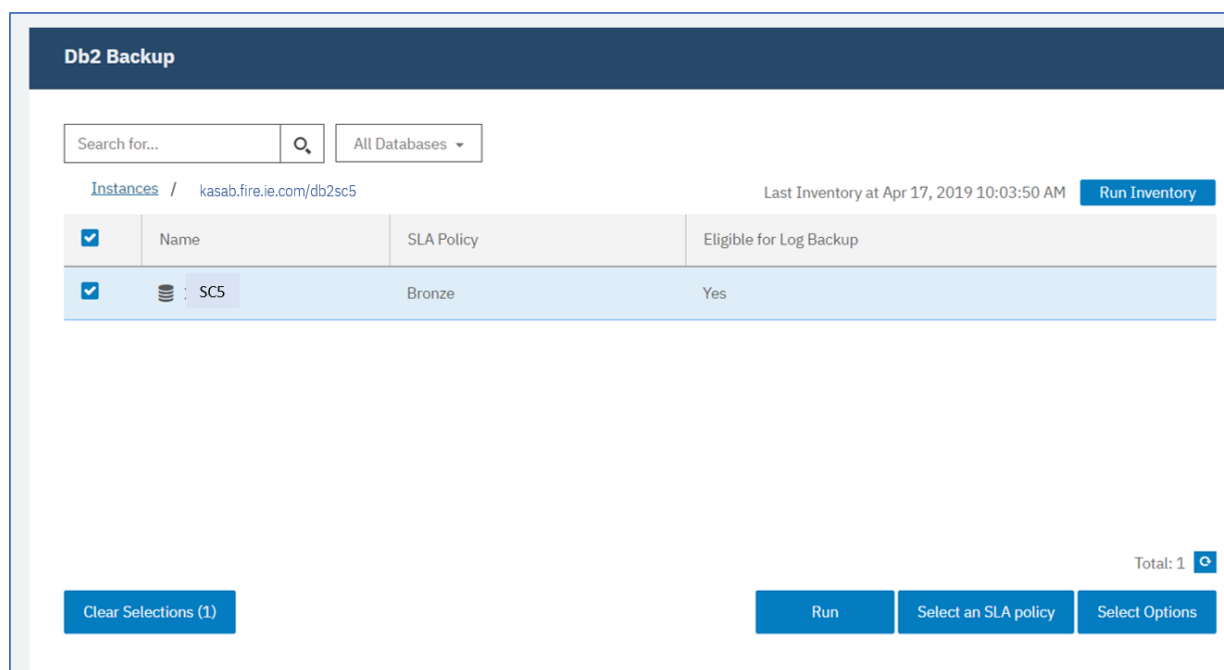
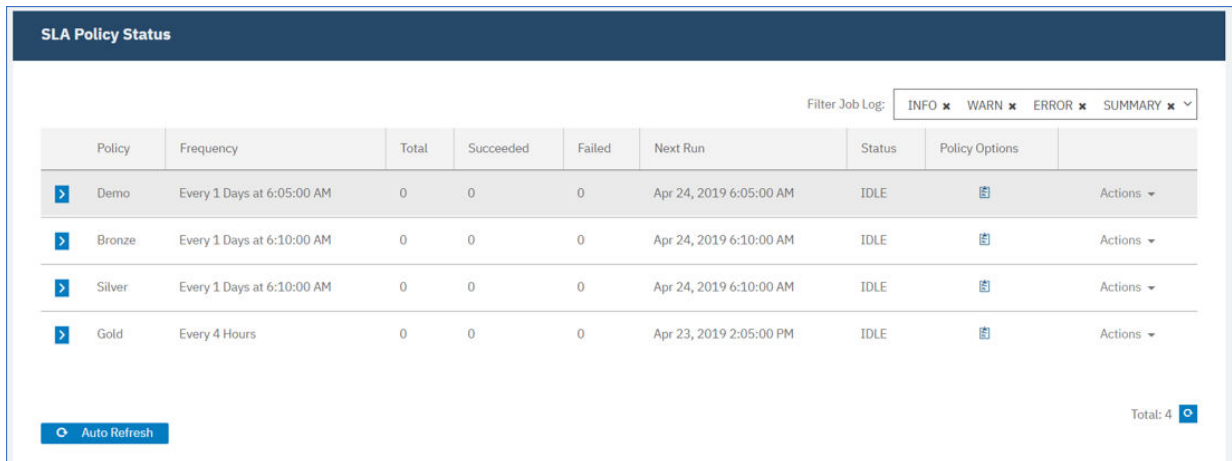


Figure 25. Db2 Backup pane showing databases in an instance

3. Click **Select SLA Policy** and select an SLA policy: **Gold**, **Silver**, or **Bronze**. Save your choice.

The predefined Gold, Silver, and Bronze policies have different frequencies and retention rates. You can create a custom SLA policy or edit an existing policy by navigating to **Policy Overview > SLA Policies**.

- 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.



Policy	Frequency	Total	Succeeded	Failed	Next Run	Status	Policy Options	Actions
Demo	Every 1 Days at 6:05:00 AM	0	0	0	Apr 24, 2019 6:05:00 AM	IDLE		Actions
Bronze	Every 1 Days at 6:10:00 AM	0	0	0	Apr 24, 2019 6:10:00 AM	IDLE		Actions
Silver	Every 1 Days at 6:10:00 AM	0	0	0	Apr 24, 2019 6:10:00 AM	IDLE		Actions
Gold	Every 4 Hours	0	0	0	Apr 23, 2019 2:05:00 PM	IDLE		Actions

Filter Job Log: INFO x WARN x ERROR x SUMMARY x

Auto Refresh

Total: 4

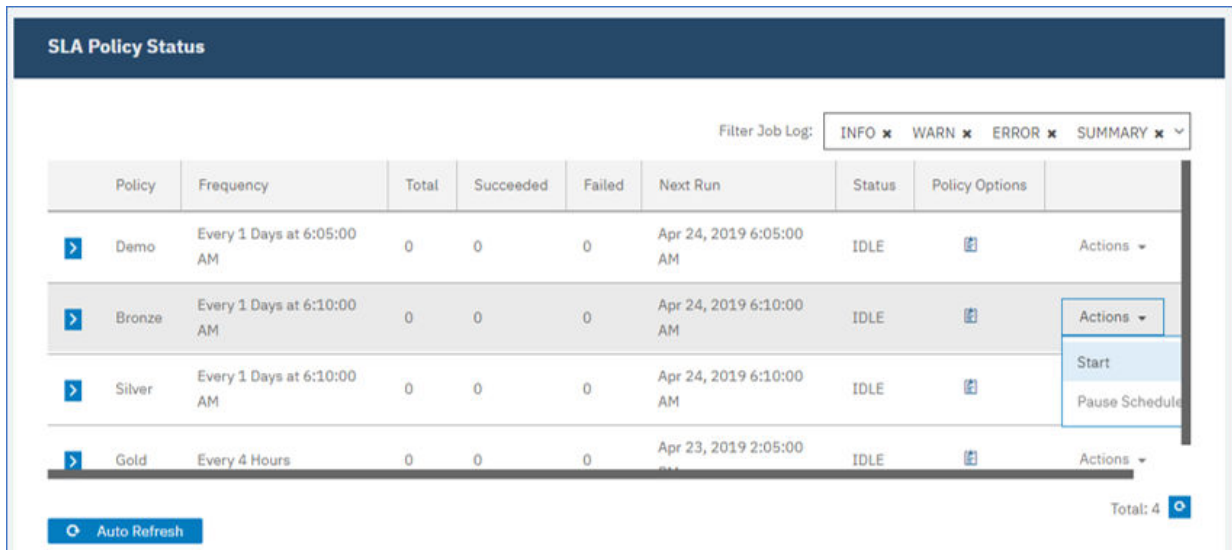
Figure 26. Backup options and SLA policies

- 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 200.

- 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 shown.



Policy	Frequency	Total	Succeeded	Failed	Next Run	Status	Policy Options	Actions
Demo	Every 1 Days at 6:05:00 AM	0	0	0	Apr 24, 2019 6:05:00 AM	IDLE		Actions
Bronze	Every 1 Days at 6:10:00 AM	0	0	0	Apr 24, 2019 6:10:00 AM	IDLE		Actions Start Pause Schedule
Silver	Every 1 Days at 6:10:00 AM	0	0	0	Apr 24, 2019 6:10:00 AM	IDLE		Actions
Gold	Every 4 Hours	0	0	0	Apr 23, 2019 2:05:00 PM	IDLE		Actions

Filter Job Log: INFO x WARN x ERROR x SUMMARY x

Auto Refresh

Total: 4

Figure 27. SLA policies

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 349.


To pause the schedule of an SLA, click **Actions** and choose **Pause Schedule**.

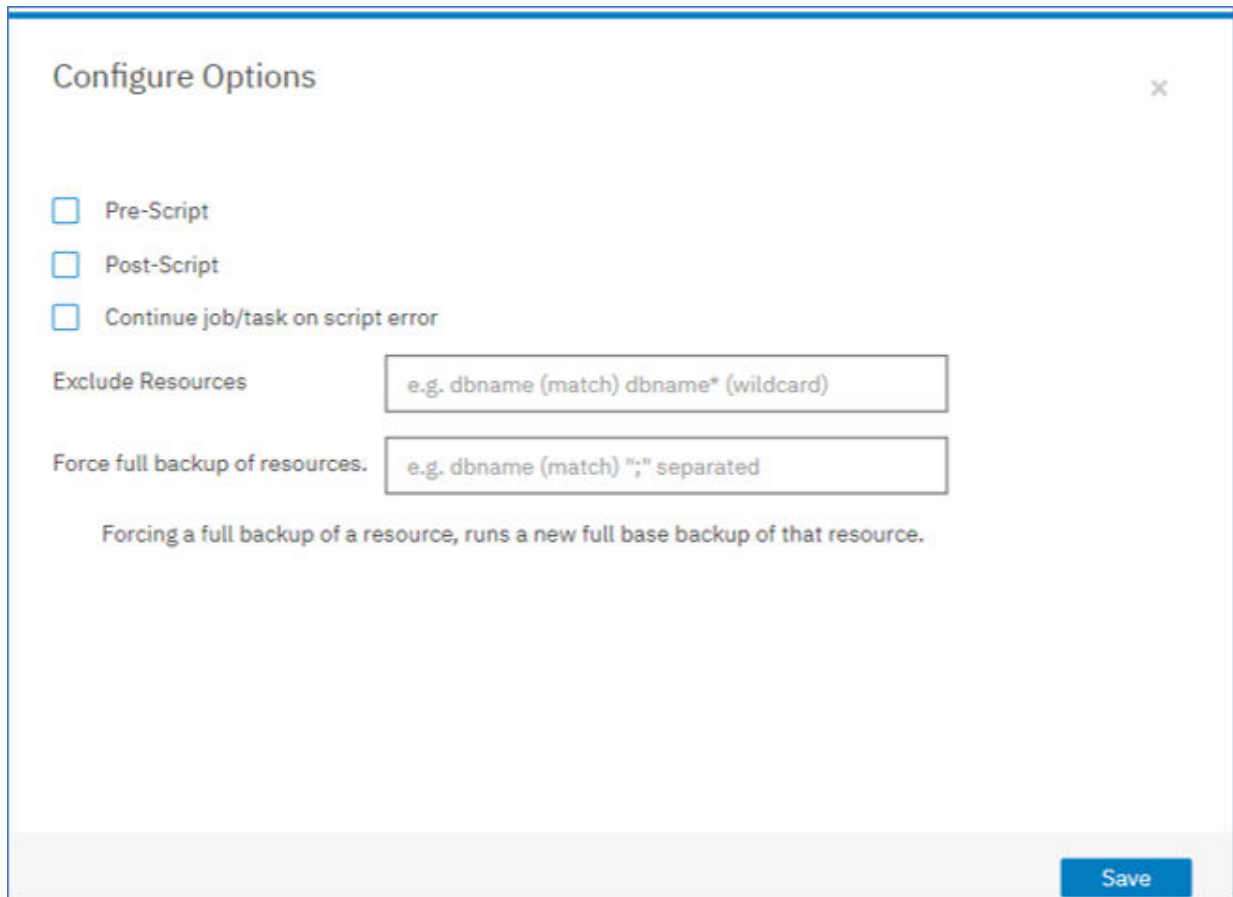
To cancel a job after it has started, click **Actions** > **Cancel**.

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.



Configure Options ×

☐ Pre-Script

☐ Post-Script

☐ Continue job/task on script error

Exclude Resources

Force full backup of resources.

Forcing a full backup of a resource, runs a new full base backup of that resource.

Save

Figure 28. Specifying SLA configuration options

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 > Applications > 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 198.

IBM Spectrum Protect Plus does not manage the retention of other archived log locations.

For more information about Db2 settings, see [IBM Db2 Welcome page](#).

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 196.
- 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 15, “Managing user access,”](#) on page 363.
- 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.

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 pane, expand **Manage Protection > Applications > Db2** and click **Create job > Snapshot restore**.

The "Snapshot restore" 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 pane 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 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> <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> • 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</p>

Option	Description
	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; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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 is selected by default to automatically clean up allocated resources as part of a restore operation when the recovery fails.
 - **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 196.
- 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 15, “Managing user access,”](#) on page 363.
- 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 pane, expand **Manage Protection > Applications > Db2** and click **Create job > Snapshot restore**.

The "Snapshot restore" 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 pane 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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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. 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 202.
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 196.
- 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 15, “Managing user access,” on page 363](#).
- 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.

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 pane, expand **Manage Protection > Applications > Db2** and click **Create job > Snapshot restore**.

The "Snapshot restore" 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 pane 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	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 that are available in your environment and are shown in the following order:

Option	Description
	<p>Backup Restores data that is backed up to a vSnap server.</p> <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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server pane.</p>
Select a location	If you are restoring data from a site, select one of the following restore locations:

Option	Description
	<p>Demo The demonstration site from which to restore snapshots.</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**. Unavailable target instances cannot be selected.


8. Choose options as described in “Restoring Db2 data ” on page 202.

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 31.

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 Administrators group and has an active Exchange mailbox in the domain.

By default, Windows adds the Exchange Organization Administrators group to other security groups, including the local Administrators group. For Exchange users who are not members of the Exchange Organization Management group, you must manually add the user account to the local Administrators group by 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.

Procedure

To add an Exchange application server, complete the following steps:

1. In the navigation pane, expand **Manage Protection > Applications > 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.
4. 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 214.
5. 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 216. For instructions about setting up Exchange database backup jobs, see [“Defining a Service Level Agreement backup job”](#) on page 218.

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 216.

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > 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 pane, expand **Manage Protection > Applications > 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 216](#).

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > 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 Operation > 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 349](#).
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 218](#).

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 pane, expand **Manage Protection > Applications > 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** and choose an SLA Policy.
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” on page 145](#).
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.
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 218](#).
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 pane, expand **Manage Protection > Applications > 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. Enter the frequency of the log backups in days, hours, or minutes.
5. Choose the start date 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.

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 pane, expand **Manage Protection > Applications > 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.
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.

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 "Snapshot 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 218](#).
- 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 15, “Managing user access,” on page 363](#).

- 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 15, “Managing user access,”](#) on page 363.

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 pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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 from the following options:

- **Test.** In test mode, the agent creates a new 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 225.](#)

6. On the **Set destination** page, select **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.

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

Enable this option to automatically clean up allocated resources as part of a restore if the recovery fails.

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 pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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.

Option	Description
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: <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; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p>

Option	Description
	<p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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.

- In the **Name** field, expand the database name to see the path information for the existing database on the application server.
- In the **New Database Name** field, enter the new name for the restored database.
- 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**.

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

Enable this option to automatically clean up allocated resources as part of a restore if the recovery fails.

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 pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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> <p>Replication Restores data that is replicated to a vSnap server.</p>

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.
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; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</p>

Option	Description
	Primary The primary site from which to restore snapshots. Secondary The secondary site from which to restore snapshots. 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	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. 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.

5. On the **Restore method** page, choose from the following options:

- **Test.** In test mode, the agent creates a new 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

Enable this option to automatically clean up allocated resources as part of a restore if the recovery fails.

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 214), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step [“15”](#) on page 236).

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 pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
 - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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


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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p>

Option	Description
	Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: Demo The demonstration site from which to restore snapshots. Primary The primary site from which to restore snapshots. Secondary The secondary site from which to restore snapshots. 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	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. 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.

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**.
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.
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 pane, click **Jobs and Operations > Active Resources** 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.
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 pane, 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.

Table 27. Previewing and filtering mailbox items	
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.

Table 27. Previewing and filtering mailbox items (continued)

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> <ol style="list-style-type: none"> 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:

Table 28. Restoring a mailbox item to another mailbox or .pst file

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 28. 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>
<p>Restore a Public Folder</p> <p> Attention:</p> <p>If a public folder mailbox is created without specifying a mailbox alias, Exchange Server uses the value of the 'DisplayName' parameter for the 'Alias' property value. If 'DisplayName' contains characters outside of unicode range from U+00A1 to U+00FF, then Alias will contain unsupported characters.</p> <p>Exchange removes unsupported characters from 'Alias' by converting them to question marks. As a result, the MMC GUI mailbox restore browser shows mailboxes with some characters displayed as '?' question marks.</p> <p>Refer to Microsoft documentation: https://docs.microsoft.com/en-us/powershell/module/exchange/mailboxes/new-mailbox?view=exchange-ps for further details.</p>	<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. In the **Jobs and Operation > Active Resources** pane, click **Actions > Cancel Granular 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 214](#)), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step [“15” on page 242](#)).

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 pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:


- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 that are 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. 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; point-in-time 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 > Site pane.</p>

Option	Description
	<p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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, 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**.
- On the **Set destination** page, select **Restore to original instance** and click **Next**.
- 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.
- 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 pane, click **Jobs and Operations > Active Resources** 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.
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 pane, 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.

Table 29. Restore options	
Option	Action
Restore Mail to Original Location	Restore mail items to their location at the time of the backup operation.

Table 29. Restore options (continued)	
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>
Restore Mail to non-Unicode PST file Restriction: <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 29. 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. In the **Jobs and Operation > Active Resources** pane, click **Actions > Cancel Granular Restore** to end the granular restore process.

Restoring Database Availability Group backups

With IBM Spectrum Protect Plus, you can restore an Exchange Server Database Availability Group (DAG) backup to the original instance or to an alternative instance.

About this task

In a DAG environment, you must restore a database to an active database copy. If you had 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.

Procedure

To define an Exchange restore job, complete the following steps:


1. In the navigation pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p>

Option	Description
	<p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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.

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

Enable this option to automatically clean up allocated resources as part of a restore if the recovery fails.

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**.

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 pane, click **Manage Protection > Applications > Exchange > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Exchange**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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> <p>Replication Restores data that is replicated to a vSnap server.</p>

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.
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; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</p>

Option	Description
	Primary The primary site from which to restore snapshots. Secondary The secondary site from which to restore snapshots. 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	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. 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.

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, go to the **Active Resources** pane, click **Actions > Cancel 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.

Ensure that your MongoDB environment meets the system requirements. For more information, see [“MongoDB requirements” on page 39](#).

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 an AIX or 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 39](#). For information about setting sudo privileges, see [“Setting sudo privileges” on page 254](#).
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 253](#).
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 when it is running on one of the following guest operating systems:

- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server Kernel-based Virtual Machine (KVM)

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 pane, expand **Manage Protection > Applications > 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. 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' interface in the MongoDB Cloud console. At the top, there's a 'Manage application servers' tab and a 'Create job' button. Below this is the 'Application Properties' section. It includes a 'Host Address' field with 'metali.ca.ibm.com' entered. There are two radio buttons: 'User' (selected) and 'SSH Key'. Under the 'User' option, there's a 'Use existing user' checkbox which is checked. Below that is a 'Select user' dropdown menu showing 'sppagent_metal.ca.ibm.com'. At the bottom of the form is a 'Get Instances' button. Below the form is a table with columns: Name, Status, Configured, and an empty column.

Figure 29. Adding a MongoDB agent

5. 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.

6. 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 15, “Managing user access,”](#) on page 363.

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 pane, expand **Manage Protection > Applications > MongoDB**.
- In the **MongoDB** window, click **Manage Application Servers**, and click **Add Application Server**.

 Add application server

- 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 254.

The Ops Manager Application Database is listed in the Instances table as shown in the following example:

```
metali8.limerick.ie.ibm.com Connection: '333.0.5.1:8888' 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 259 and [“Defining a regular service level agreement job”](#) on page 260.

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

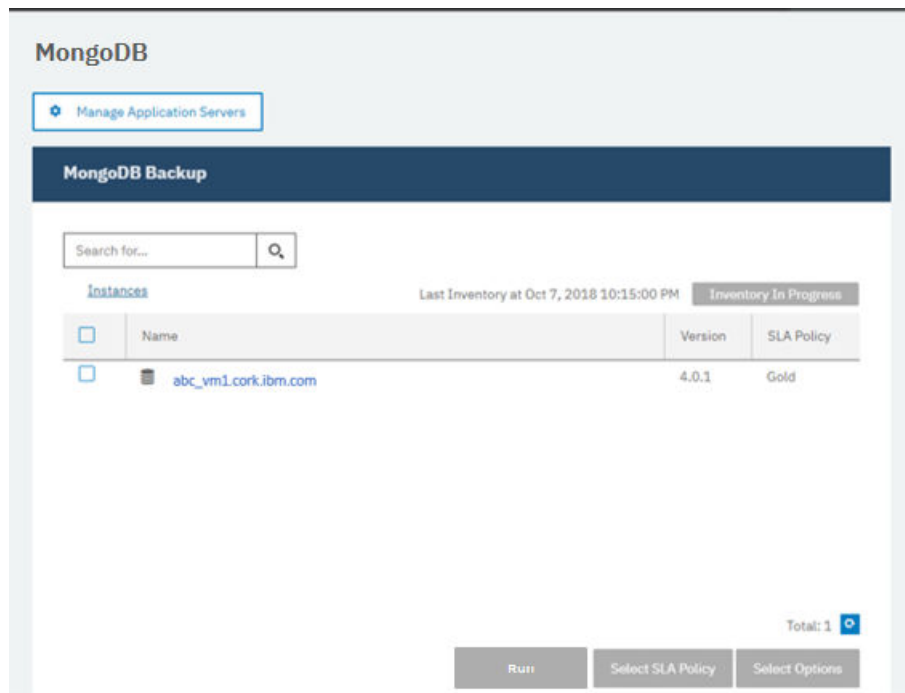
Ensure that you added your MongoDB application servers to IBM Spectrum Protect Plus. For instructions, see [Adding a MongoDB application server](#).

Procedure

- In the navigation pane, expand **Manage Protection > Applications > 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 pane, click **Manage Protection > Applications > 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.

1. Physical - Basic Test for physical host network configuration

Name	Description	Status	Message
Host FQDN Resolvable Test	Host FQDN must be resolvable to an IPv4 address	✓	
Socket Connection Test	Must allow socket connection on port 22 for Linux	✓	

2. Remote - Remote executor test for session creation and remote agent deployment

Name	Description	Status	Message
Remote Session Test	Latest remote agent must be installed on host, SSH and SFTP service must be installed on Linux host, and port must be open to create session to SSH service.	✓	
Remote Agent Execute Test	Remote agent must be configured correctly using user credentials with sufficient privileges.	✓	

3. LINUX - Basic Linux prerequisites for file and volume operations

Name	Description	Status	Message
Sudo Privileges	User must have password-less sudo privileges	✓	

OK

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 15, “Managing user access,”](#) on page 363 and [“Roles for MongoDB”](#) on page 253.

Restriction: Do not run inventory jobs at the same time that backup jobs are scheduled.

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 349.

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > 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 Operation > 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 349.
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.

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 pane, expand **Manage Protection > Applications > MongoDB**.
2. Select the MongoDB instance to back up all the data in that instance.

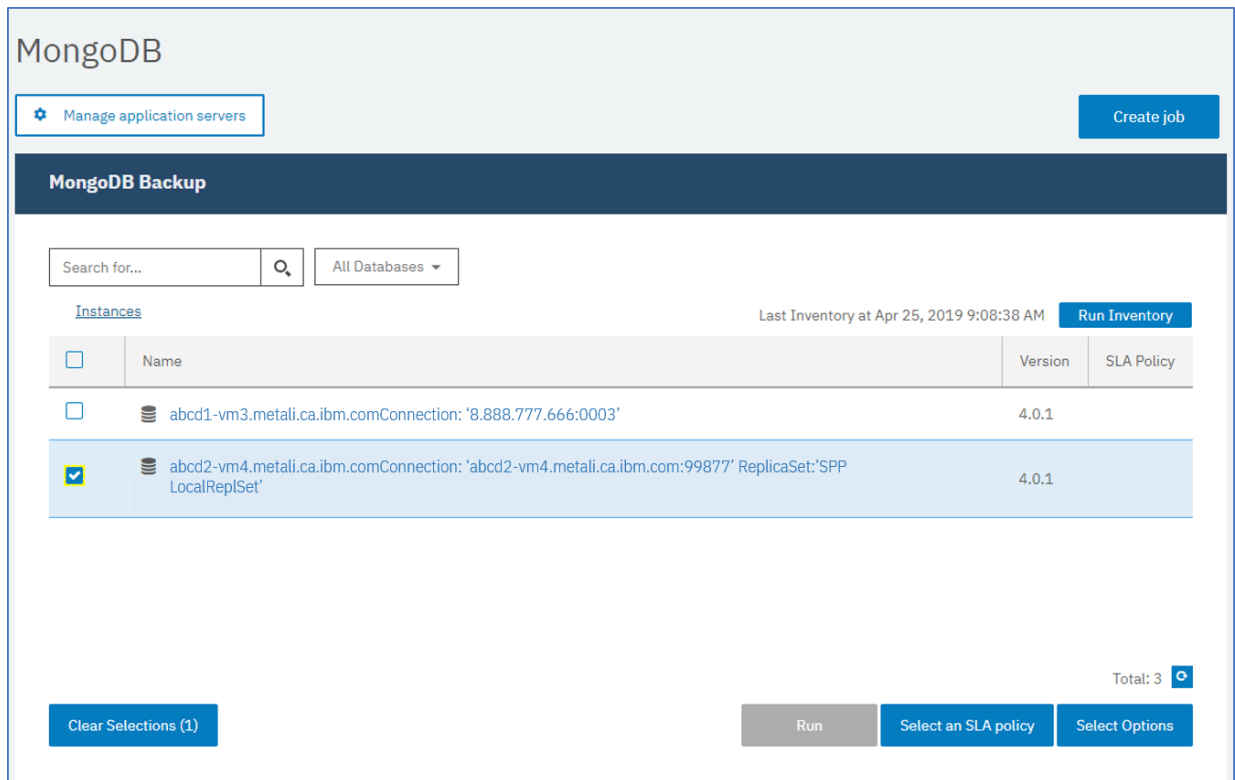


Figure 30. MongoDB Backup pane showing instances

3. Click **Select an SLA policy** and choose an SLA policy. Save your choice.

Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. You can also create a custom SLA policy by navigating to **Policy Overview > Add SLA Policy**.

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.

Options

Maximum Parallel Streams per Database:

SLA Policy Status

Filter Job Log: Info x Warning x Error x Summary x v

Policy	Frequency	Total	Succeeded	Failed	Next Run	Status	Policy Options	Actions
Gold	Every 4 Hours	1	1	0	Apr 25, 2019 10:05:00 AM	Idle		

Figure 31. Backup options and SLA Policy Status

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 261.

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.

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.

Configure Options ×

☐ Pre-Script

☐ Post-Script

☐ Continue job/task on script error

Exclude Resources

Force full backup of resources.

Forcing a full backup of a resource, runs a new full base backup of that resource.

Save

Figure 32. Specifying additional SLA configuration options

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 259](#).
- 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 15, “Managing user access,” on page 363](#), and [“Roles for MongoDB” on page 253](#).
- 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 pane, click **Manage Protection > Applications > MongoDB > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > MongoDB**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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 is selected by default to automatically clean up allocated resources as part of a restore operation if recovery fails.

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 Snapshot 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**.

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 259.
- 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 15, “Managing user access,”](#) on page 363, and [“Roles for MongoDB”](#) on page 253.
- 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 pane, click **Manage Protection > Applications > MongoDB > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > MongoDB**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p>

Option	Description
	<p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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.

- **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 is selected by default to automatically clean up allocated resources as part of a restore operation if recovery fails.

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.

8. 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 Snapshot 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**.

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 259](#).
- 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 15, “Managing user access,” on page 363](#), and [“Roles for MongoDB” on page 253](#).
- 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 pane, click **Manage Protection > Applications > MongoDB > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > MongoDB**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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> <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>

Option	Description
	<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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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>

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 is selected by default to automatically clean up allocated resources as part of a restore operation if recovery fails.

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.

8. 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 Snapshot 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**.

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 pane, click **Manage Protection > Applications > MongoDB > Create job**, and then select **Snapshot restore** to open the "Snapshot 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**.
12. Click **Actions > Cancel Restore** to end the granular restore procedure.

Microsoft Office 365

To protect Microsoft™ Office 365 email, calendars, contacts, and data on OneDrive cloud storage, you must first register the Office 365 application with Azure Active Directory. Then, deploy the application server and register it with IBM Spectrum Protect™ Plus. After that, you must add Office 365 tenants, and define a service level agreement (SLA) policy to create backup jobs.

Registering with Azure Active Directory

To protect an Office 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 Office 365 subscription.
- Ensure that you have an Office 365 administrative user ID and password.

Procedure

1. Go to the Office 365 welcome page and sign in to your Microsoft account by using your Office 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 Office 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 Office 365 application with IBM Spectrum Protect Plus.
8. To create a client secret for this application ID, click **Certificates & secrets > New client secret**.
9. On the "Add a client secret" pane, enter any user name in the **Description** field, and click **Add**. A client secret is generated, and the value is then 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. Select the API name, for example, Azure Active Directory Graph.
 - a) For permission name User.Read.All, select the **Delegated Permissions** type.
 - b) For the remaining permissions, select the **Application Permissions** type for each permission name for the API in the table.

API	Permission name
Azure Active Directory Graph	User.Read.All
Azure Active Directory Graph	Directory.Read.All
Exchange	full_access_as_app
Microsoft Graph	Calendars.ReadWrite
Microsoft Graph	Contacts.ReadWrite
Microsoft Graph	Files.ReadWrite.All
Microsoft Graph	Mail.ReadWrite
Microsoft Graph	Sites.Read.All
Microsoft Graph	User.Read
Microsoft Graph	User.Read.all

13. To save the selected permissions, click **Grant admin consent for <your organization name>**.

What to do next

Follow the instructions in [“Registering the Office 365 tenant with IBM Spectrum Protect Plus”](#) on page 279.

Registering the Office 365 tenant with IBM Spectrum Protect Plus

To ensure that the IBM Spectrum Protect™ Plus agent can connect to the Office 365 tenant, you must register the Office 365 tenant credentials, and the proxy host server with IBM Spectrum Protect Plus. This procedure is necessary to ensure that Office 365 data can be backed up to IBM Spectrum Protect Plus.

Before you begin

Ensure that the Office 365 application is registered with Azure Active Directory. For instructions, see [“Registering with Azure Active Directory”](#) on page 278.

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > Office 365**.

2. On the Office 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: This is the Organization/Tenant name such as *tenantname.onmicrosoft.com*. This is NOT the user facing name that you set up (which appears under **Display name**) when doing Azure app registration.
 - 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 host name or IP of the Linux server that is being used as the proxy host.
 - b. 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 in IBM Spectrum Protect Plus, an inventory is run automatically on the Office 365 organization, which returns the Office 365 users in that resource.

Backing up Office 365 data

After your Office 365 organization is registered with IBM Spectrum Protect Plus, you can apply a service level agreement (SLA) policy to start protecting the Office 365 data.

Procedure

1. In the IBM Spectrum Protect Plus navigation pane, expand **Manage Protection > Applications > Office 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 93.
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.
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, 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 your chosen SLA and you can follow the progress of the job in the log.

Incremental forever backup for Office 365

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.

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
- 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 Office 365 accounts.
2. All subsequent backup jobs back up only new or changed data from the selected accounts.

Restoring Office 365 data

You can restore Office 365 data from backup copies on vSnap servers or remote storage. When you are ready to restore a mailbox to Office 365, you can complete the task in IBM Spectrum Protect™ Plus.

Before you begin

At least one Office 365 backup job must have run successfully. For instructions about setting up a backup job, see [“Backing up Office 365 data”](#) on page 280.



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 pane, expand **Manage Protection > Applications > Office 365**.
2. Click **Create job**.
3. Select **Snapshot 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. Then, click **Next** to continue.
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 Office 365 account Restore to another account Restores data to another Office 365 account
Restore Path	Restores data to selected directory path in the Office 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. Then, click **Next** to continue.
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 pane, click **Jobs and Operations** and then click the **Running Jobs** tab.

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 44.

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 pane, click **Manage Protection > Applications > Oracle > Backup**.
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. Select **User** or **SSH key**.

Option	Description
User	<p>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:</p> <p>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.</p> <p>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.</p> <p>For Kerberos-based authentication only, the user identity must be specified in the <i>username@FQDN</i> format. The user name must be able to authenticate using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain that is specified by the fully qualified domain name.</p> <p>Password Enter your password for the application server.</p>
SSH Key	Click this option to use an SSH key. Select a key from the Select a SSH key list.

6. 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.
7. 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.
8. 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 368 .

Related concepts

[“Managing user access” on page 363](#)

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 284](#)

Use a backup job to back up Oracle environments with snapshots.

[“Restoring Oracle data” on page 287](#)

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 File 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.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Applications > Oracle > Backup**.
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 pane, click **Manage Protection > Applications > Oracle > Backup**.
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.

- 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 15, “Managing user access,”](#) on page 363.
- Register the providers that you want to back up. For more information, see [“Adding an Oracle application server”](#) on page 282.
- Configure SLA policies. For more information, see [“Create backup policies”](#) on page 93.

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 282.

Procedure

To define an Oracle backup job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Applications > 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 349.
5. To edit options before you create the job definition, click **Select Options**. Set the job definition 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.

Select **Truncate source logs after successful backup** to automatically delete older archived logs from the database's primary archived log location. If the option is cleared, 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. If the option is selected, IBM Spectrum Protect Plus deletes older unneeded archived logs from the primary log location at the end of every successful database backup.

When the option **Truncate source logs after successful backup** is selected, set the retention of primary logs through the **Primary log retention in days** setting. This setting controls the quantity of archived logs that are retained in the primary archived log locations. For example, if **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.

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**.
7. 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 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 287 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 350](#)

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 284](#).
- 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 15, “Managing user access,” on page 363](#).

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 pane, click **Manage Protection > Applications > Oracle > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.

Tips:

- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > Oracle**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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 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 resources restore; point-in-time 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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p>

Option	Description
	Repository server archive The repository server to which snapshots were copied to tape. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: Demo The demonstration site from which to restore snapshots. Primary The primary site from which to restore snapshots. Secondary The secondary site from which to restore snapshots. 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	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. 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.

- 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.

- 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.

- 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.

This option is applicable only when you are restoring an Oracle database to its original location by using its original database name.

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

Enable this option to automatically clean up allocated resources as part of a restore operation if the recovery fails.

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.

8. 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.

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.

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 350](#)

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 282](#)

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 49](#).

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

Kerberos requirements

Kerberos-based authentication can be enabled through a configuration file on the IBM Spectrum Protect Plus appliance. This will override the default Windows NTLM protocol.

For Kerberos-based authentication only, the user identity must be specified in the username@FQDN format. The username must be able to authenticate using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain specified by the fully qualified domain name.

Kerberos authentication also requires that the clock skew between the Domain Controller and the IBM Spectrum Protect Plus appliance is less than five minutes.

The default Windows NTLM protocol is not time dependent.

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 SQL Server instance can use a specific user account to access the resources of that particular 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: <https://support.microsoft.com/en-us/help/968264/error-message-when-you-try-to-map-to-a-network-drive-of-a-dfs-share-by>

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

To add an SQL Server host, complete the following steps.

1. In the navigation pane, click **Manage Protection > Applications > SQL > Backup**.
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.

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.

For Kerberos-based authentication only, the user identity must be specified in the *username@FQDN* format. The user name must be able to authenticate using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain that is specified by the fully qualified domain name.

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 368 .

Related concepts

[“Managing user access” on page 363](#)

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 296](#)

Use a backup job to back up SQL Server environments with snapshots.

[“Restoring SQL Server data” on page 300](#)

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 pane, click **Manage Protection > Applications > SQL > Backup**.
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 pane, click **Manage Protection > Applications > SQL > Backup**.
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 15, “Managing user access,” on page 363](#).
- 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](#).

- SQL databases from test restores are not eligible for backup operations. As a result, SQL databases that are a product of a test restore cannot be selected for backup nor will they be selectable in an SLA. Finally, if an SLA association is at the instance level, any SQL databases from test restores will be skipped during backup operations.
- 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.
- For more information about log backups for SQL, see [“Log backups” on page 299](#).

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 294](#).
- Configure SLA policies. For more information, see [“Create backup policies” on page 93](#).
- 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 pane, click **Manage Protection > Applications > 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**.
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 349](#).
5. Click **Select Options** to specify more options before you save the backup job.

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.

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 more 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. 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**.

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 300 .

Related concepts

“Configuring scripts for backup and restore operations” on page 350

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 345

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.

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, follow the instructions in [“Backing up SQL Server data” on page 296](#).

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).
- SQL databases that result from test restore operations are not eligible for backup operations. As a result, SQL databases that are a product of a test restore cannot be selected for backup nor will they be selectable in an SLA. Finally, if an SLA association is at the instance level, any SQL databases from test restores are skipped during backup operations.
- 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 296.
- 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 15, “Managing user access,” on page 363.
- If you are planning to run a point-in-time recovery, ensure that both the restore target SQL instance 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 <https://docs.microsoft.com/en-us/sql/t-sql/language-reference?view=sql-server-2017>.

- 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.
- SQL databases restored from test restores are not eligible for backups and log backups.

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 share. 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.

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 SQL restore job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Applications > SQL > Create job**, and then select **Snapshot restore** to open the "Snapshot restore" wizard.


Tips:


- You can also open the wizard by clicking **Jobs and Operations > Create job > Snapshot restore > SQL**.
- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation pane 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	<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: <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; point-in-time restore; or recurring restore

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 > Site pane.</p> <p>Cloud service The cloud service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object Storage pane.</p> <p>Repository server The repository server to which snapshots were copied. The repository server is defined in the System Configuration > Backup Storage > Repository Server pane.</p> <p>Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the System Configuration > Backup Storage > Object 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 > Backup Storage > Repository Server 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.</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.

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.

6. 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.
- **By Transaction ID.** Select this option to configure a point-in-time recovery by transaction ID.

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

Automatically clean up allocated resources as part of a restore operation if the recovery fails.

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.

8. 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.

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.

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 350](#)

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 294](#)

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 296](#)

Use a backup job to back up SQL Server environments with snapshots.

Chapter 11. Protecting containers

Kubernetes Backup Support is a feature of IBM Spectrum Protect Plus that extends data protection to containers in Kubernetes clusters. Kubernetes is a system for orchestrating containers across clusters of hosts.

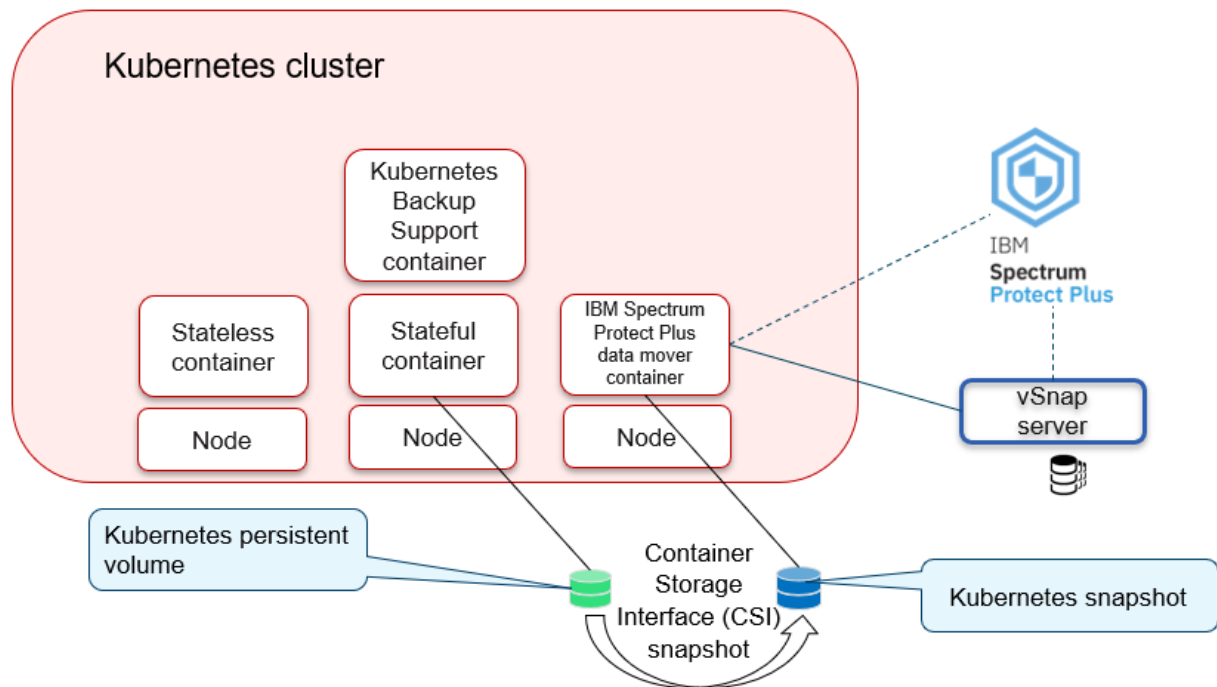
Overview of Kubernetes Backup Support

IBM Spectrum Protect Plus Kubernetes Backup Support protects persistent volumes that are attached to containers in Kubernetes clusters. Snapshot backups of the persistent volumes are created and copied to IBM Spectrum Protect Plus vSnap servers.

Persistent volumes that contain application data are protected by predefined service level agreement (SLA) policies that specify how often snapshot and copy backups are created and how long they are retained. If data on the original volumes is damaged or lost, the volumes can be restored from either the snapshot or copy backups on the vSnap servers.

Kubernetes Backup Support protects only persistent storage that was allocated by a storage plug-in that supports the Container Storage Interface (CSI) provided for Kubernetes. Kubernetes Backup Support is fully tested with Red Hat Ceph block storage, which supports CSI. The CSI plug-in provides snapshot capabilities that are used for backup operations.

The following figure shows how Kubernetes Backup Support is deployed in the Kubernetes environment and how it interacts with IBM Spectrum Protect Plus:



Data mover container

The IBM Spectrum Protect Plus data mover is deployed in a container in the Kubernetes environment. The data mover container communicates with the IBM Spectrum Protect Plus instance outside of the Kubernetes environment for copy backup support.

Kubernetes Backup Support uses persistent volume claims (PVCs) to identify the persistent volumes to back up. When a backup schedule is run, snapshot and copy backups of a PVC are created at the time intervals that are specified by the SLA. The data mover copies the data and records the snapshot backups

in the IBM Spectrum Protect Plus **Jobs and Operations** window. Snapshots that are created by on-demand backups are also recorded in IBM Spectrum Protect Plus.

Multitenancy is supported

Kubernetes Backup Support manages backup and restore operations by using Kubernetes custom resources. All backup and restore objects belong to a Kubernetes namespace. The Kubernetes administrator can restrict access to these objects. With controlled access, multiple users can run backup and restore requests in the same Kubernetes cluster. The backup and restore objects inherit a namespace from the PVC that identifies the persistent volume for backup and restore operations. For more information about multitenancy, see [“Security features in Kubernetes Backup Support” on page 311](#).

Backup and restore types

Kubernetes Backup Support provides multiple types of backup and restore functions. Backup and restore operations are initiated by Kubernetes requests for services.

Backup types

The following types of backup operations are available:

Snapshot backup

Creates a backup of the persistent volume by using Container Storage Interface (CSI) storage plug-in snapshot capabilities. The snapshot is stored in a location that is assigned by a Kubernetes snapshot class as defined by the backup administrator. Typically, this location is the same storage site as the persistent volume that is being backed up. The snapshot class must be compatible with the storage class of the persistent volume. In other words, the snapshot class and storage class are defined and provided by the same CSI storage plug-in.

Snapshot backups are created by scheduled backup requests and on-demand backup requests. On-demand backup requests are available only on volumes that are already protected by scheduled backups.

During scheduled backups, snapshot and copy backups are created at intervals that are defined by a service level agreement (SLA) policy. During an on-demand backup request, a snapshot is taken immediately but no copy backup is created.

Copy backup

Copies the full persistent volume to an IBM Spectrum Protect Plus vSnap server. Based on predefined SLA policies, IBM Spectrum Protect Plus offers longer retention of copy backups compared to snapshot backups.

During scheduled backups, snapshot and copy backups are created at intervals that are defined by the SLA policy.

Restriction: You cannot create copy backups of raw volumes because unformatted volumes cannot be mounted to the data mover container for copy backup operations.

Restore types

The following types of restore operations are available:

Snapshot restore

Restores a snapshot to a new persistent volume. This type of operation is suitable for rapidly restoring recent snapshot backups.

Copy backup restore

Restores a copy backup to the original persistent volume or to a new persistent volume. If you want to restore a copy backup to the original persistent volume, the container to which the persistent volume is attached must not be running.

This type of operation is suitable for restoring persistent volumes from copy backups that are retained for a longer period on IBM Spectrum Protect Plus. For example, you can use this type of operation to restore copy backups whose snapshots have expired.

SLA policies

Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained.

The following predefined service level agreement (SLA) policies are available to help you protect your persistent volumes:

Table 30. SLA policies for Kubernetes Backup Support		
SLA policy	Snapshot backup frequency and retention period	Copy backup frequency and retention period
test	15-minute intervals and 1 hour	Hourly and 1 day
daily	4-hour intervals and 24 hours	Daily and 31 days
weekly	Daily and 7 days	Weekly and 31 days
monthly	None	31-day intervals and 365 days

SLA policies are predefined and cannot be modified. You can associate only one SLA with a volume. The SLA is assigned to a volume in the scheduled backup definition.

When snapshots expire, they are deleted automatically in the Kubernetes environment. When copy backups expire, they are marked for expiration on IBM Spectrum Protect Plus and are deleted by IBM Spectrum Protect Plus maintenance jobs.

The SLA policies are available in the ConfigMap object that is named `baas-sla` in the `baas` namespace. To view this `baas-sla` ConfigMap, issue the following command:

```
kubect1 describe configmap baas-sla -n baas
```

Restriction: In a production environment, do not schedule backups with the **test** SLA policy. Backups with the **test** SLA in a production environment are not supported. The **test** SLA is provided as a means for you to test your setup for Kubernetes Backup Support. Use the **test** SLA only to schedule backup jobs of small sample volumes for testing purposes. Entries are added to the log to indicate that the backups were created for testing. After you validate that the scheduled backup jobs ran correctly and you can successfully create snapshot and copy backups in IBM Spectrum Protect Plus, discontinue using the **test** SLA.

User roles

Depending on their role, enterprise developers and backup administrators interact with different user interfaces to protect persistent data in containers.

Enterprise developer

The enterprise developer uses the Kubernetes command-line tool (**kubect1**) to complete the following tasks independent of the backup administrator:

- Initiates self-service backup and restore requests
- Selects a service level agreement (SLA) policy to use in backup requests to protect their volumes
- Views the status of backup and restore requests
- Queries information about snapshot and copy backups
- Pauses and resumes scheduled backups operations
- Removes obsolete scheduled backup requests and on-demand snapshot requests

Backup administrator

The backup administrator completes the following tasks:

- Deploys and sets up Kubernetes Backup Support software in the Kubernetes environment
- Creates the Kubernetes storage class for persistent volumes and the snapshot class for storing snapshots
- Installs and configures IBM Spectrum Protect Plus
- Monitors copy backup jobs by using the IBM Spectrum Protect Plus user interface
- Generates reports that show the history of container backup jobs by using the IBM Spectrum Protect Plus user interface
- Completes troubleshooting tasks, such as collecting log files for debugging in the Kubernetes environment and viewing trace log files for Kubernetes Backup Support

Kubernetes Backup Support requests

To protect your container data, you submit Kubernetes Backup Support requests in the Kubernetes environment.

A Kubernetes Backup Support request is a Kubernetes custom resource that is of kind BaaSReq. The requests are specified in *YAML Ain't Markup Language* (YAML) configuration files. The request is then submitted by using the **kubect1** command-line interface.

Types of requests in Kubernetes Backup Support

The following table shows the available types of Kubernetes Backup Support requests. The request types are specified as values for the **requesttype** key in the YAML file.

Table 31. Types of Kubernetes Backup Support requests	
Request type	Description
Backup	Schedule a backup operation for a PVC
Restore	Restore a PVC from a snapshot backup or a copy backup
Pause	Pause a scheduled backup for a PVC
Resume	Resume a paused scheduled backup
Destroy	Delete all snapshot and copy backups and mark the scheduled job as destroyed
OnDemandBackup	Request an immediate snapshot backup of the PVC

Running a request

To initiate a request, create a YAML configuration file that specifies the request type and provide the required parameters. Then, submit the request by running the **kubect1 create** command.

The following sample file (baas-req.yaml) shows the general format of a YAML file:

```
#-----  
# Filename: baas-req.yaml  
#-----  
  
apiVersion: "baas.io/v1alpha1"  
kind: BaaSReq  
  
metadata:  
  name: request_name  
  namespace: namespace  
spec:
```

```
requesttype: request_type  
sla: test | daily | weekly | monthly
```

where:

request_name

The name of the request. For scheduled backup, pause, and resume requests, the name of the request must match the PVC name.

namespace

The namespace in which the persistent volume exists. If you do not specify a namespace, the default namespace is used.

request_type

The type of request. For the list of available request types, see [“Types of requests in Kubernetes Backup Support”](#) on page 310.

test | daily | weekly | monthly

A predefined service level agreement (SLA) policy that you can assign to the request. For pause and resume requests, this key is ignored. For more information, see [“SLA policies”](#) on page 309.

To start the request that is specified in the `baas-req.yaml` sample file, enter the following command on the command line:

```
kubectl create -f baas-req.yaml
```

To check the status of a request, use one of the following methods:

- To list all Kubernetes Backup Support requests in all namespaces that you can access, enter the following command:

```
kubectl get baasreq --all-namespaces
```

- To display the status of all Kubernetes Backup Support requests in the specified namespace, enter the following command:

```
kubectl describe baasreq -n namespace
```

where *namespace* is the namespace in which the persistent volume exists.

- To display the status of a specific Kubernetes Backup Support request, enter the following command:

```
kubectl describe baasreq request_name -n namespace
```

where *request_name* is the name of the request, and *namespace* is the namespace in which the persistent volume exists.

Security features in Kubernetes Backup Support

In addition to basic security features that are integrated into Kubernetes Backup Support, advanced security features are provided to help protect containers, secure network connections, encrypt data, and verify installation packages.

Security scanning of containers

Kubernetes Backup Support components are built on containers that are derived from the Red Hat Universal Based Image (UBI). The Kubernetes Backup Support software on each container was statically scanned for vulnerable components or libraries. In addition, the containers are dynamically scanned to help prevent runtime vulnerabilities such as code injection. After the scan, the software is tested by using an automated test suite to verify that Kubernetes Backup Support can operate as expected and correctly process erroneous input.

All containers, except for the data mover container, run in a dedicated namespace that provides further security isolation. The data mover must run in the same namespace as the persistent volume claim (PVC) for backup or restore operations because the mounting of the volume is limited to containers in a single namespace.

Least privileged containers

Each of the components in Kubernetes Backup Support runs under the principle of least privilege. The actions of the containers are constrained by the role-based authentication control rules that are associated with their service accounts in their separate namespace. In addition, the software in each container runs as a non-root user. Only the data mover runs as a privileged container because the data mover requires access to the mount point on the host system of the volume that is backed up or restored. All other containers are not privileged.

Authentication of network connections

The network connections between Kubernetes Backup Support components are controlled by network policies that limit the connections to the ones that are required for correct operation. Connections to IBM Spectrum Protect Plus rely upon the security protocols that are provided by IBM Spectrum Protect Plus.

Multitenancy

Multitenancy is supported in Kubernetes Backup Support, which relies extensively on the authentication and authorization that is provided by the Kubernetes cluster for namespaces. Because the authorization is related to a namespace, any user who is authorized to create a BaaSReq object in that namespace can request a backup or restore for any PVC that is associated with that namespace. A BaaSReq object is a custom Kubernetes resource that is used in Kubernetes Backup Support requests.

Snapshots are protected by the Container Storage Interface (CSI) to restrict access to the namespace of the original PVC. Kubernetes Backup Support associates the namespace with the backup copies that are stored in IBM Spectrum Protect Plus, and the backup copies must be restored to volumes in the same namespace.

Encryption of data at rest

The cluster and storage administrators are responsible for enabling the mechanisms for protecting data at rest through encryption. The sensitive data includes the copy backup data and Kubernetes Backup Support secrets, which consist of user IDs and passwords that were specified during the installation process. The cluster administrator can specify that secrets are encrypted when stored in the cluster etcd database. For more information, see [Encrypting Secret Data at Rest](#).

Kubernetes Backup Support does not implement additional encryption beyond what is provided by the cluster. However, the storage administrator can deploy an IBM Spectrum Protect Plus vSnap server that is enabled for encryption. When enterprise developers create requests to back up PVCs, the developers can specify encryption as part of the backup requests.

When encryption is specified, the backup request is sent to the IBM Spectrum Protect Plus server. When the request is received by the IBM Spectrum Protect Plus server, it directs the data to a vSnap server for encryption if the vSnap server is enabled for encryption of data at rest. If the backup request specified encryption but encryption is not enabled on the vSnap server for the data at rest, no error is issued. In this case, the data is not encrypted.

The IBM Spectrum Protect Plus can also confirm whether the data is in an encrypted vSnap server.

Code signing

The cluster administrator can verify that the Kubernetes Backup Support installation package has not been modified since it was generated by IBM. This process is accomplished by verifying the signature file that is included with the installation package against the appropriate signature and certificates. The verification process is described in the installation documentation.

For more information, see [“Installing and deploying Kubernetes Backup Support images”](#) on page 315.

Installing Kubernetes Backup Support

The backup administrator must install and configure Kubernetes Backup Support in the Kubernetes environment.

Prerequisites for Kubernetes Backup Support

Before you can install Kubernetes Backup Support, ensure that all system requirements and prerequisites are met.

For Kubernetes Backup Support system requirements, see [“Kubernetes Backup Support requirements” on page 54](#).

Then, to meet the prerequisites for Kubernetes Backup Support, complete the following actions in the Kubernetes environment:

- [“Enabling the VolumeSnapshotDataSource feature” on page 313](#)
- [“Verifying whether the metrics server is running” on page 314](#)
- [“Defining the application and persistent volume claim relationship” on page 314](#)

Enabling the VolumeSnapshotDataSource feature

To support copy backup and snapshot restore operations, you must enable the **VolumeSnapshotDataSource** alpha feature.

For more information about alpha features, see [Feature Gates](#).

To enable the **VolumeSnapshotDataSource** alpha feature, you must patch the Kubernetes scheduler, controller, and API server as follows:

1. Using the **sudo** command, edit the following YAML files:

```
/etc/kubernetes/manifests/kube-apiserver.yaml
/etc/kubernetes/manifests/kube-controller-manager.yaml
/etc/kubernetes/manifests/kube-scheduler.yaml
```

2. In each YAML file, add the following statement within the command section:

```
- --feature-gates=VolumeSnapshotDataSource=true
```

Important: Ensure that you edit the YAML files directly and do not create backup copies of these files in the same directory. The presence of the backup copies in the `/etc/kubernetes/manifests` directory might negate the changes that you made to enable the **VolumeSnapshotDataSource** feature gate.

You might have to wait a minute or two for the changes to be detected by Kubernetes.

3. Verify whether the feature is enabled by issuing the following commands:

```
ps aux | grep apiserver | grep feature-gates
```

```
ps aux | grep scheduler | grep feature-gates
```

```
ps aux | grep controller-manager | grep feature-gates
```

The output for one of these commands is similar to the following example:

```
root      13121  7.4  2.5 518276 305424 ?        Ssl  Sep06 120:37 kube-apiserver --
authorization-mode=Node,RBAC --advertise-address=192.0.2.0
--allow-privileged=true --client-ca-file=/etc/kubernetes/pki/ca.crt --enable-admission-
plugins=NodeRestriction --enable-bootstrap-token-auth=true
--etcd-cafile=/etc/kubernetes/pki/etcd/ca.crt --etcd-certfile=/etc/kubernetes/pki/apiserver-
etcd-client.crt --etcd-keyfile=/etc/kubernetes/pki/apiserver-etcd-client.key
--etcd-servers=https://127.0.0.1:2379 --insecure-port=0 --kubelet-client-certificate=/etc/
kubernetes/pki/apiserver-kubelet-client.crt
```

```
--kubelet-client-key=/etc/kubernetes/pki/apiserver-kubelet-client.key --kubelet-preferred-
address-types=InternalIP,ExternalIP,Hostname
--proxy-client-cert-file=/etc/kubernetes/pki/front-proxy-client.crt --proxy-client-key-
file=/etc/kubernetes/pki/front-proxy-client.key
--requestheader-allowed-names=front-proxy-client --requestheader-client-ca-file=/etc/
kubernetes/pki/front-proxy-ca.crt
--requestheader-extra-headers-prefix=X-Remote-Extra- --requestheader-group-headers=X-Remote-
Group --requestheader-username-headers=X-Remote-User
--secure-port=6443 --service-account-key-file=/etc/kubernetes/pki/sa.pub --service-cluster-
ip-range=198.51.100.0/24 --tls-cert-file=/etc/kubernetes/pki/apiserver.crt
--tls-private-key-file=/etc/kubernetes/pki/apiserver.key --feature-
gates=VolumeSnapshotDataSource=true
```

Verifying whether the metrics server is running

To help optimize product performance, ensure that Kubernetes Metrics Server 0.3.5 or later is installed and running properly on your cluster. The metrics server is required for the Kubernetes Backup Support scheduler to determine the resources that are used by concurrent data mover instances.

If the metrics server does not return data, the number of data movers that are used for backup operations is limited, which might negatively impact performance.

You can verify that the metrics server is installed and returning metrics data by completing the following steps:

1. Verify the installation by issuing the following command:

```
kubectl get deploy,svc -n kube-system | egrep metrics-server
```

The output is similar to the following example:

```
deployment.extensions/metrics-server 1/1 1 1 3d4h
service/metrics-server ClusterIP 198.51.100.0 <none> 443/TCP 3d4h
```

2. Verify that the metrics server is returning data for all nodes by issuing the following command:

```
kubectl get --raw "/apis/metrics.k8s.io/v1beta1/nodes"
```

The output is similar to the following example:

```
{"kind": "NodeMetricsList", "apiVersion": "metrics.k8s.io/v1beta1", "metadata": {"selfLink": "/
apis/metrics.k8s.io/v1beta1/nodes"}, "items": [{"metadata":
{"name": "cirrus12", "selfLink": "/apis/metrics.k8s.io/v1beta1/nodes/cirrus12",
"creationTimestamp": "2019-08-08T23:59:49Z"}, "timestamp": "2019-08-08T23:59:08Z",
"window": "30s", "usage": {"cpu": "1738876098n", "memory": "8406880Ki"}}]}
```

Tip: The command might fail with empty output for the "items" key. This error is likely caused by installing the metrics server with a self-signed certificate. To resolve this issue, install the metrics server with a correctly signed certificate that is recognized by the cluster.

Defining the application and persistent volume claim relationship

You can optionally tie your stateful applications to their persistent volume claims (PVCs) by using an owner-dependent relationship. By defining this relationship, you enable cascading actions for the applications.

For example, scaling up and scaling down an application can cause the scheduled backups of its PVC to be paused and resumed. Similarly, deleting the application causes the deletion of the PVC, which in turn triggers the deletion of the backups.

After an application starts using a PVC to store persistent data, you can reconfigure the PVC definition with its owner application.

The following example is a sample configuration file for a PVC that shows the owner-dependent relationship between an application and a PVC object. The PVC object includes the details of the owner deployment.

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: demo-pvc
  ownerReferences:
    - apiVersion: apps/v1beta1
      blockOwnerDeletion: true
      kind: Deployment
      name: Dept10-deployment
      uid: 3b760e89-7da5-11e9-8c5a-0050568ba59c
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
  storageClassName: csi-rbd
```

Installing and deploying Kubernetes Backup Support images

Before you can back up and restore persistent volumes that are attached to your containers, you must install and deploy Kubernetes Backup Support images.

Before you begin

Complete the following tasks:

- Ensure that your system environment meets the requirements that are described in “Kubernetes Backup Support requirements” on page 54 and “Prerequisites for Kubernetes Backup Support” on page 313.
- Download the installation file `installer-10.1.5.tar.gz` from Passport Advantage® Online. For information about downloading files, see [technote 1072392](#).
- Validate the downloaded file by using one of the following methods:
 - Verify the MD5 checksum of the downloaded installation file. Ensure that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
 - Verify the signed file that is associated with the installation package by issuing the following command:

```
openssl dgst -sha256 -verify IBMSPSignCertificatePublic -signature ./
installer-10.1.5.tar.gz.sig ./installer-10.1.5.tar.gz
```

About this task

During the installation and deployment procedure, you must first update the `baas_config.cfg` configuration file with specifications for your environment, and then run the installation script `baas_install.sh`. When you run the installation script, an appropriate Helm Chart is automatically called to deploy Kubernetes Backup Support in your environment.

Procedure

Complete the following steps on the command line in the Kubernetes environment:

1. Log in to the operating system on the master node of the Kubernetes cluster that is used as the installation node.
2. Unpack the installation package (`installer-10.1.5.tar.gz`) by entering the following command:

```
tar -xvf installer-10.1.5.tar.gz
```

This command extracts a folder that is named `installer`.

3. Go to the installer directory by entering the following command:

```
cd installer
```

4. Run the following two commands to obtain the Classless Inter-Domain Routing (CIDR) method for the cluster and the IP address and port for the cluster API server. The values are used in Step “5” on page 316.

- a) Obtain the CIDR for the cluster by issuing the following command:

```
kubectl cluster-info dump | grep -m 1 cluster-cidr
```

The CIDR is provided in the output in the following format:

```
--cluster-cidr=xxx.yyy.0.0/zz
```

The CIDR is similar to the following example:

```
198.51.0.0/24
```

- b) Obtain the IP address and server port for the cluster API server by issuing the following command:

```
kubectl config view|awk '/cluster\:\/\/,server\:\/\/' | grep server\: | awk '{print $2}'
```

The result is a URL that is composed of an IP address and port number, as shown in the following example:

```
https://192.0.2.0:6443
```

where 192.0.2.0 is the cluster API server IP address and 6443 is the port address.

5. Edit the `baas_config.cfg` file with a text editor and modify the configuration parameters by providing the appropriate values for your environment. Enclose the values in quotation marks, as shown in the following example.

```
BAAS_ADMIN="sppadmin"
```

For parameters that contain a list of values, provide the list of values in a comma-separated format that is enclosed in quotation marks, as shown in the following example:

```
SPP_VSNAP_IP_ADDRESSES="192.0.2.0,192.0.2.1"
```

The following table contains the parameters that you must modify:

Table 32. Specifications for the <code>baas_config.cfg</code> configuration file	
Parameter	Description
BAAS_ADMIN	The user ID of the IBM Spectrum Protect Plus administrator.
BAAS_PASSWORD	The IBM Spectrum Protect Plus password. For increased security, specify an empty string (""). You are prompted for the password when you run the deployment script. If you must specify a password in the configuration file for automated test deployments, ensure that the file is stored in a secure location.

Table 32. Specifications for the <code>baas_config.cfg</code> configuration file (continued)	
Parameter	Description
DATAMOVER_USER	<p>The IBM Spectrum Protect Plus application host user name.</p> <p>You can use the default data mover name or specify a different name. This user account is automatically configured and used in the data mover container.</p>
DATAMOVER_PASSWORD	<p>The IBM Spectrum Protect Plus application host password.</p> <p>For increased security, specify an empty string (" "). You are prompted for the password when you run the deployment script. If you must specify a password in the configuration file for automated test deployments, ensure that the file is stored in a secure location.</p>
CLUSTER_CIDR	The CIDR for the cluster. Enter the CIDR that was obtained in Step “4.a” on page 316.
CLUSTER_API_SERVER_IP_ADDRESS	The IP address for the cluster API server. Enter the IP address that was obtained in Step “4.b” on page 316.
CLUSTER_API_SERVER_PORT	The port address for the cluster API server. Enter the port address that was obtained in Step “4.b” on page 316.
SPP_IP_ADDRESSES	The IBM Spectrum Protect Plus server IP address.
SPP_VSNAP_IP_ADDRESSES	<p>The IP address for the IBM Spectrum Protect Plus vSnap server.</p> <p>You can obtain this address from the IBM Spectrum Protect Plus user interface by clicking System Configuration > Backup Storage > Disk > Disk Storage. This parameter can contain more than one IP address. Provide the list of values in a comma-separated format that is enclosed in quotation marks, as shown in the following example:</p> <pre>SPP_VSNAP_IP_ADDRESSES="192.0.2.0,192.0.2.1"</pre>
PRODUCT_IMAGE_REGISTRY	<p>The Docker registry address and port that hosts the containers.</p> <p>Enter the address in the <code>ip_address:port</code> format.</p>
PRODUCT_IMAGE_REGISTRY_NAMESPACE	The Docker registry namespace that hosts the containers.

Table 32. Specifications for the `baas_config.cfg` configuration file (continued)

Parameter	Description
PRODUCT_IMAGE_REGISTRY_SECRET_NAME	<p>The name of the Kubernetes image-pull secret that contains the credentials for the registry. The secret must be in the namespace that is specified by the <code>PRODUCT_IMAGE_REGISTRY_NAMESPACE</code> parameter.</p> <p>If you are using an internal registry, enter an empty string ("").</p> <p>For the data mover container to run, the image-pull secret must be in every namespace of each persistent volume claim (PVC) to be backed up and restored.</p>

Restrictions:

- The following parameters and values are reserved for Kubernetes Backup Support. Keep them as is.

```
PRODUCT_NAMESPACE="baas"
PRODUCT_TARGET_PLATFORM="K8S"
```

- The `SPP_PORT` value specifies the port for the Kubernetes Backup Support user interface. Do not change the default value of 443.
- Kubernetes Backup Support is available only in English in IBM Spectrum Protect Plus Version 10.1.5. For this reason, do not change the `PRODUCT_LOCALIZATION="en_US"` setting.

Your specifications are automatically inserted into the ConfigMap (`baas-configmap`) during the deployment.

- Start the installation and deployment by issuing the following command.

```
./baas_install.sh -i
```

All container images are in the image registry and are running.

When prompted, enter yes to continue.

A project namespace for the Kubernetes Backup Support deployment called "baas" is created. This project is created before the images are pushed into the image registry, which is identified by the namespace.

Depending on your environment, it might take several minutes to load and deploy the package.

- To verify that the Kubernetes Backup Support components are properly installed, issue the following command:

```
./baas_install.sh -s
```

If the installation fails, the missing components are listed in the `MISSING` section of the output.

Tip: You can also check the status of the installation with the `./helm status baas` command.

Results

When all pods are running, the deployment is completed. To verify that all pods are in the Running state and no components are missing, issue the following command:

```
kubectl get pods
```

or

```
kubectl describe pod pod_name
```

For example, you can issue the following command:

```
kubectl get pods -n baas
```

The output is similar to the following example:

NAME	READY	STATUS	RESTARTS	AGE
baas-controller-59dbcf7f94-c6zjr	1/1	Running	0	3h56m
baas-datamover-b44f755c5-k5g2f	1/1	Running	0	22h
baas-etcd-client-59bd5d647f-k76b	1/1	Running	0	2d21h
baas-scheduler-55944fbbb6-b96lw	1/1	Running	0	3h56m
baas-transaction-manager-856b7fd6c94h558	1/1	Running	0	3h32m
baas-etcd-spp-job-control-store-65d9dfb84d-vlqb5	1/1	Running	2	2d17h

If the data mover container is not listed in the output, the data mover container is deployed at run time.

You can show the Kubernetes Backup Support services that are set up by issuing the following command:

```
kubectl get services -n baas
```

The output is similar to the following example:

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
baas-etcd-client	ClusterIP	10.103.44.178	<none>	2379/TCP	2d21h
baas-etcd-spp-job-control-store	ClusterIP	10.100.229.67	<none>	2379/TCP	2d21h
baas-scheduler	ClusterIP	10.96.33.79	<none>	8000/TCP	2d21h
baas-transaction-manager	ClusterIP	10.106.230.11	<none>	5000/TCP	2d21h

The `baas-datamover` service is deployed at runtime with type `NodePort` instead of the `ClusterIP` range with the TCP protocol.

You can show the Kubernetes Backup Support network policies that are deployed by issuing the following command:

```
kubectl get networkpolicies -n baas
```

The output is similar to the following example:

NAME	POD-SELECTOR	AGE
baas-ctl-networkpolicy	app=baas,component=controller,release=baas	2d21h
baas-etcd-networkpolicy	app=baas,component=etcd-client,release=baas	2d21h
baas-etcd-spp-job-control-store	app=baas,component=etcd-spp-job-control-store,release=baas	2d21h
baas-scheduler	app=baas,component=scheduler,release=baas	2d21h
baas-transaction-manager	app=baas,component=transaction-manager,release=baas	2d21h

The network policy for the data mover is deployed at runtime with the pod-selector `app=baas,component=datamover,release=baas`.

What to do next

After the deployment is completed, ensure that backup operations run correctly by using the test service level agreement (SLA) policy to run scheduled backups. For instructions, see [“Scheduling backups of persistent volumes”](#) on page 322.

If you want to update the existing configuration or to upgrade an existing installation of Kubernetes Backup Support, modify the parameters in the `baas_config.cfg` file as required for your environment, and issue the following command:

```
./baas_install.sh -u
```

Related concepts

[“Troubleshooting Kubernetes Backup Support”](#) on page 377

To help troubleshoot issues with Kubernetes Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Uninstalling Kubernetes Backup Support

You can uninstall Kubernetes Backup Support while keeping the customizations that were applied to the namespace. By keeping the customizations, you make it easier to redeploy Kubernetes Backup Support in the future if necessary.

Before you begin

Take the following actions before you begin the uninstallation:

- Pause all scheduled backups. For instructions, see [“Pausing scheduled backups” on page 333](#).
- Wait for all running backup and restore jobs to finish.
- If you do not want to keep the snapshots on the persistent volume after the uninstallation, delete all your snapshot backups before you begin the uninstallation. For instructions, see [“Deleting container backups” on page 334](#).

About this task

After you uninstall Kubernetes Backup Support, you can remove the Kubernetes Backup Support images to save space on your file system.

Procedure

To uninstall Kubernetes Backup Support from the Kubernetes cluster that you are logged in to, complete the following steps on the command line:

1. Go to the `installer` directory by issuing the following command:

```
cd installer
```

2. To uninstall Kubernetes Backup Support, issue the following command:

```
./baas_install.sh -d
```

3. When prompted, enter yes to continue.

The Kubernetes Backup Support deployment is removed from the Kubernetes environment.

4. Optional: To verify the progress of the uninstallation, enter the following command:

```
kubectl get pod -n baas
```

5. Optional: To remove all Kubernetes Backup Support images from your Kubernetes environment, issue the following command:

```
docker images -a | grep "baas" | awk '{print $3}' | xargs docker rmi -f
```

Results

The `baas_install.sh` installation script does not automatically delete the Kubernetes Backup Support product namespace ("baas") that is specified in the `baas_config.cfg` file. The customizations that were applied to the namespace are preserved, so that you can reuse the namespace for the reinstallation of Kubernetes Backup Support. If you want to delete the namespace, you must do it manually.

What to do next

If you want to fully uninstall Kubernetes Backup Support, see [“Completely uninstalling Kubernetes Backup Support” on page 321](#).

Completely uninstalling Kubernetes Backup Support

You can completely uninstall Kubernetes Backup Support so that all components, including all configurations and backups, are removed from the Kubernetes environment. After you remove the components, you can also remove Kubernetes Backup Support images to save space on your file system.

Before you begin

Take the following actions before you begin the uninstallation:

- Pause all scheduled backups. For instructions, see [“Pausing scheduled backups” on page 333](#).
- Wait for all running backup and restore jobs to finish.

Procedure

To completely uninstall Kubernetes Backup Support from the cluster that you are logged in to, complete the following steps on the command line:

1. Destroy all snapshot and copy backups with a **destroy** request. For instructions, see [“Deleting container backups” on page 334](#).
2. Delete any persistent volume claims (PVCs) that were used for copy backups.

Tip: You can look for the names of the PVCs that were backed up.

3. Delete the baas custom resource definition (CRD) by issuing the following command:

```
kubectl delete crd baasreqs.baas.io
```

This command also deletes all BaasReq request objects.

4. Uninstall Kubernetes Backup Support by issuing the following command from the `installer` directory:

```
./baas_install.sh -d
```

When prompted, enter yes to continue.

This command removes all data mover pods, deployments, and network policies.

5. Optional: To verify the progress of the uninstallation, enter the following command:

```
kubectl get pod -n baas
```

6. Disable the **VolumeSnapshotDataSource** feature if you no longer require it.
7. Delete the service level agreement (SLA) policies and any other customizations by deleting the baas namespace. Issue the following command:

```
kubectl delete namespace baas
```

8. Optional: Review the installation and configuration information and revert any prerequisite steps.
9. To remove all Kubernetes Backup Support images from your Kubernetes environment, issue the following command:

```
docker images -a | grep "baas" | awk '{print $3}' | xargs docker rmi -f
```

Backing up container data

To protect persistent volumes that are attached to a container, you can schedule backup operations to run as specified by predefined service level agreements (SLAs). You can also create snapshots of your persistent volumes immediately by running on-demand backup requests.

Scheduling backups of persistent volumes

You can create scheduled backup requests that run based on predefined service level agreement (SLA) policies that specify how often backup operations are run and how long snapshot and copy backups are retained.

Before you begin

Backup requests are directed to persistent volume claims (PVCs) for the volumes that you want to protect. Ensure that the PVC exists within the specified namespace.

The PVC must be formatted for it to be backed up. For the PVC be formatted correctly, it must be mounted and written to. Backup operations of raw block volumes are not supported.

About this task

When a scheduled backup job runs, a snapshot of the persistent volume is created and a snapshot of the volume is copied to an IBM Spectrum Protect Plus vSnap server at the frequency that is defined by the SLA. For example, for the **daily** SLA policy, a snapshot is taken every 4 hours and a snapshot is copied to the vSnap server every 24 hours.

All backup jobs are scheduled, except for on-demand backup jobs. To schedule backup jobs for a PVC, create a YAML configuration file with specifications for the backup job and apply the request on the Kubernetes command line.

Procedure

1. Optional: Display a list of PVCs in your namespace by issuing the following command:

```
kubectl get pvc -n namespace
```

From the list of PVCs, identify the PVC that you want to back up.

2. Create a YAML file that defines the request for a scheduled backup. The YAML file must contain the following properties:

```
#-----  
# Filename: filename.yaml  
#-----  
  
apiVersion: "baas.io/v1alpha1"  
kind: BaaSReq  
  
metadata:  
  name: pvc_name  
  namespace: namespace  
spec:  
  requesttype: Backup  
  sla: test | daily | weekly | monthly  
  encryption: no | yes  
  volumesnapshotclass: snapshot_class_name
```

where:

filename

The name of the YAML configuration file. The file type is `.yaml`.

pvc_name

The name of the backup request, which must match the name of the PVC for the volume that you want to back up. For example, if you want to set up the backup request for the PVC named dbvol-01, the name of the request must also be dbvol-01.

namespace

The namespace in which the PVC exists.

sla: test | daily | weekly | monthly

The SLA policy that determines the schedule for backup operations. Specify one of the following SLAs:

Table 33. SLA policies for Kubernetes Backup Support		
SLA policy	Snapshot backup frequency and retention period	Copy backup frequency and retention period
test	15-minute intervals and 1 hour	Hourly and 1 day
daily	4-hour intervals and 24 hours	Daily and 31 days
weekly	Daily and 7 days	Weekly and 31 days
monthly	None	31-day intervals and 365 days

Restriction: In a production environment, do not schedule backups with the **test** SLA policy. Backups with the **test** SLA in a production environment are not supported. The **test** SLA is provided as a means for you to test your setup for Kubernetes Backup Support. Use the **test** SLA only to schedule backup jobs of small sample volumes for testing purposes. Entries are added to the log to indicate that the backups were created for testing. After you validate that the scheduled backup jobs ran correctly and you can successfully create snapshot and copy backups in IBM Spectrum Protect Plus, discontinue using the **test** SLA.

If you want to change the values for the **sla** or other parameters in the YAML file and apply it to the same PVC, see [Modifying parameters in a YAML file](#).

encryption: no | yes

Specify whether to encrypt the copy backup data that is stored on IBM Spectrum Protect Plus. For encryption to occur, the IBM Spectrum Protect Plus vSnap server must be enabled for encryption. For more information, see [“Encryption of data at rest” on page 312](#).

Specify one of the following values:

no

Do not encrypt the copy backup data on IBM Spectrum Protect Plus. This value is the default.

yes

Encrypt the copy backup data on IBM Spectrum Protect Plus. If you specified **yes** but the vSnap server is not set up for encryption, the data is not encrypted and no error is indicated.

snapshot_class_name

The snapshot class for the snapshot backups. This value is optional. If you do not specify the snapshot class, the default snapshot class is used if the snapshotter in the default snapshot class matches the provisioner of the volume. Otherwise, the backup request is invalid.

3. Start the backup schedule by applying the backup request. Enter the following command on the command line:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

After you submit the backup request, the first scheduled backup operation will start within the window that is defined by the SLA policy. The time of the backup is recorded in the backup status.

Kubernetes Backup Support takes ownership of all snapshots, regardless of how they were created.

What to do next

To view information about the backup, issue the **kubectl describe** command by using the request name or the PVC name. For more information, see [“Viewing the status of backup and restore jobs” on page 329](#).

Modifying parameters in a YAML file: You cannot change the **sla** or any other parameter after a scheduled backup has started on a PVC. Simply updating parameters in a YAML file and then applying the YAML file will not change the values in a running backup request.

However, if you must update the **sla** or other parameters after the schedule has started on a PVC, complete the following steps:

1. Delete your container backups with a **destroy** request. For instructions about deleting backups, see [“Deleting container backups” on page 334](#).
2. Check the status of the **destroy** request and ensure that it is finished. To check the status of the request, issue the following command:

```
kubectl describe baasreq pvc_name -n namespace | grep Backupstatus
```

In the command output, ensure that the **Backupstatus** field is shown as follows:

```
Backupstatus: Destroyed
```

3. Delete the original scheduled backup request with the following command:

```
kubectl delete baasreq pvc_name -n namespace
```

4. Update the parameters in the YAML file.
5. Create a new scheduled backup with the updated YAML file:

```
kubectl create -f filename.yaml
```

Related concepts

[“Backup and restore types” on page 308](#)

Kubernetes Backup Support provides multiple types of backup and restore functions. Backup and restore operations are initiated by Kubernetes requests for services.

[“SLA policies” on page 309](#)

Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained.

[“Kubernetes Backup Support requests” on page 310](#)

To protect your container data, you submit Kubernetes Backup Support requests in the Kubernetes environment.

[“Troubleshooting Kubernetes Backup Support” on page 377](#)

To help troubleshoot issues with Kubernetes Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Backing up a persistent volume on demand

If you do not want to wait for a scheduled backup job to run or if you must create a snapshot immediately, you can run an on-demand backup job. You can run an on-demand backup job only if the volume is already protected by scheduled backup jobs.

Before you begin

Backup requests are directed to persistent volume claims (PVCs) for the volumes that you want to protect. Ensure that the PVC exists within the specified namespace.

The PVC must be formatted for it to be backed up. For the PVC be formatted correctly, it must be mounted and written to. Backup operations of raw block volumes are not supported.

You can run an on-demand backup job for a persistent volume only if the volume is protected by scheduled backup jobs. For more information, see [“Scheduling backups of persistent volumes” on page 322](#).

About this task

During an on-demand backup operation, only a snapshot backup is created.

Unlike the request for scheduled backups, the name of the on-demand request must be unique. In other words, the name of the request must not be the same as the name of the PVC.

Procedure

1. Optional: Display a list of PVCs in your namespace by issuing the following command:

```
kubectl get pvc -n namespace
```

From the list of PVCs, identify the PVC that you want to back up.

2. Create a YAML file that defines the request for an on-demand backup operation. The YAML file must contain the following properties:

```
#-----  
# Filename: filename.yaml  
#-----  
  
apiVersion: "baas.io/v1alpha1"  
kind: BaaSReq  
  
metadata:  
  name: name_of_request  
  namespace: namespace  
spec:  
  requesttype: OnDemandBackup  
  pvcname: pvc_name
```

where:

filename

The name of the YAML configuration file. The file type is `.yaml`.

name_of_request

The name of the on-demand backup request. The name must be unique, and must not match the name of the PVC.

A new on-demand backup request must be created for each subsequent on-demand backup of the same PVC. In other words, if you want to create a second on-demand backup of a PVC, create a new request and specify a different request name (*name_of_request*) in the YAML file.

namespace

The namespace in which the PVC exists.

pvc_name

The name of the PVC for the volume that you want to back up.

3. Start the on-demand backup by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

Kubernetes Backup Support takes ownership of all snapshots, regardless of how they were created.

Snapshots are retained according to the retention period that is specified by the service level agreement (SLA) policy that is associated with the volume. When a snapshot is expired based on the SLA policy of the

PVC, the snapshot is deleted. The request for the on-demand backup job is updated to show that the snapshot has expired, as shown by the **Backupstatus** field.

To view information about the backup, issue the **kubect1 describe** command by using the request name or the PVC name. For more information, see [“Viewing the status of backup and restore jobs” on page 329](#).

Related concepts

[“Backup and restore types” on page 308](#)

Kubernetes Backup Support provides multiple types of backup and restore functions. Backup and restore operations are initiated by Kubernetes requests for services.

[“Kubernetes Backup Support requests” on page 310](#)

To protect your container data, you submit Kubernetes Backup Support requests in the Kubernetes environment.

[“Troubleshooting Kubernetes Backup Support” on page 377](#)

To help troubleshoot issues with Kubernetes Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Restoring container data

You can restore a persistent volume from a snapshot backup or a copy backup. A snapshot restore operation is generally the faster method for restoring a persistent volume.

Before you begin

For any type of restore, you cannot restore a volume to a different namespace.

You can restore a snapshot only to a new persistent volume. If you are restoring a snapshot, the persistent volume claim (PVC) for the new volume is automatically created when you restore the snapshot.

You can restore a copy backup to a new or original persistent volume. If you are restoring a copy backup to a new persistent volume, the PVC for the new volume is automatically created when you restore the copy backup. Except for the original location, the restore fails if you specify a PVC that already exists.

If you are restoring a copy backup to the original persistent volume, the application container to which the persistent volume is attached must not be running.

Restriction: To ensure that a restore request works correctly, do not manually delete any snapshots of volumes that are protected by Kubernetes Backup Support.

About this task

Depending on your recovery point objective and recovery time objective, you can run a fast restore, a copy restore, or a fast -ondemand restore operation.

The following scenarios can help you select the type of restore operation:

- To rapidly restore a recent snapshot that was created as part of a schedule, run a fast restore operation. If another operation is in progress on the same volume, the fast restore operation might take longer to complete.
- To restore a volume from a particular point in time after the corresponding snapshot has expired, run a copy restore operation to restore the copy backup from IBM Spectrum Protect Plus.
- To restore a snapshot from an on-demand backup, run a fast -ondemand restore.
- To verify a copy backup before it is restored to the original volume, you can run a copy restore operation to restore the copy backup to a new volume. Then, you can verify the contents of the new volume. If no issues are found in the new volume, you can restore the copy backup to the original volume.

Restore points are identified by the time stamp of the snapshot or copy backup.

Procedure

1. To show the restore points that are available for a PVC, query all the backups for the PVC by running the following command:

```
kubectl describe BaaSReq pvc_name -n namespace
```

2. In the status output that is displayed, identify the time stamp of the source snapshot or copy backup that you want to restore. The time stamps are shown in the Status section of the output before the type of backup.

For example, the following output shows the time stamps for different types of backups:

```
Status:
Timestamp: 2019-05-30 13:27:21
Type:      FAST
Timestamp: 2019-05-30 13:32:21
Type:      COPY
Timestamp: 2019-06-11 18:59:46
Type:      FAST-ONDEMAND
```

where:

FAST

Denotes the backup type for a snapshot backup that is taken during a scheduled backup operation.

COPY

Denotes the backup type for a copy backup that is stored on an IBM Spectrum Protect Plus vSnap server.

FAST-ONDEMAND

Denotes the backup type for an on-demand snapshot backup.

3. Create a YAML file for the restore request that contains the following properties. Insert the time stamp for the source snapshot in the **restorepoint** parameter.

```
#-----
# Filename: filename.yaml
#-----

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: name_of_restore_request
  namespace: namespace
spec:
  requesttype: restore
  pvcname: pvc_name
  targetvolume: target_volume_for_restore
  storageclass: storage_class_of_target_volume
  restorepoint: timestamp_of_backup
  restoretype: fast | copy | fast-ondemand
```

where:

filename

The name of the YAML configuration file.

name_of_restore_request

The name of the request for the restore job. The name must be unique, and must not be the same as the name of the PVC.

A new restore request must be created for each subsequent restore of the same PVC. In other words, if you want to restore a PVC again, create a new request and specify a different request name (*name_of_request*) in the YAML file.

namespace

The namespace for the request.

pvc_name

The name of the PVC that you want to restore.

target_volume_for_restore

The name of the PVC that you want to restore the volume to.

For fast restores, the volume is always restored to a new PVC. In this case, provide the name of the new PVC.

For copy restores, you can restore the volume to the original or new PVC. If you are restoring a copy backup to a new persistent volume, the PVC for the new volume is automatically created when you restore the copy backup. Except for the original location, the restore fails if you specify a PVC that already exists.

storage_class_of_target_volume

The storage class that is defined for the target volume.

For fast restore operations, the storage class is ignored. The storage class of the original PVC is used.

For copy restore operations, you can specify a storage class that is the same as the original PVC or specify a different storage class. If you do not specify the storage class, the storage class of the original PVC is used.

If you specify a storage class but do not specify the restore type with the **restoretype** parameter, a copy restore is performed.

timestamp_of_backup

The time stamp of the source snapshot or copy backup that you want to restore from. The time stamp is in Coordinated Universal Time (UTC) format.

If you do not specify a time stamp, the most recent snapshot or copy backup is restored.

restoretype: fast | copy | fast-ondemand

The type of restore operation to use.

fast

Restores a volume from a snapshot backup that was created as part of a scheduled backup.

copy

Restores a volume from a copy backup.

fast-ondemand

Restores a volume from an on-demand snapshot backup.

This parameter is optional. If you do not specify a restore type, the type of restore is determined automatically. If a snapshot exists at the specified time stamp, a fast restore is run to restore the snapshot. If only a copy backup is available at the specified time, a copy restore is run to restore the copy backup.

4. Start the restore request by entering the following command on the command line:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

What to do next

If you restored data to a new persistent volume, you can reconfigure the application container to mount the new volume after the snapshot or copy backup is restored.

As a best practice, delete the completed request by issuing the following command:

```
kubectl delete baasreq name_of_restore_request -n namespace
```

Deleting completed requests has the following benefits:

- It reduces the size of the etcd database and allows you to reuse the name of a request for another operation.
- It makes troubleshooting easier.

- It makes it easier for you to track backup and restore requests that are running in your Kubernetes cluster.
- At any point in time, you have a clear picture of requests that are running in on your cluster when you issue the following command:

```
kubectl get baasreq -n namespace
```

Related concepts

[“Backup and restore types” on page 308](#)

Kubernetes Backup Support provides multiple types of backup and restore functions. Backup and restore operations are initiated by Kubernetes requests for services.

[“Kubernetes Backup Support requests” on page 310](#)

To protect your container data, you submit Kubernetes Backup Support requests in the Kubernetes environment.

[“Troubleshooting Kubernetes Backup Support” on page 377](#)

To help troubleshoot issues with Kubernetes Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Related tasks

[“Viewing the status of backup and restore jobs” on page 329](#)

After you submit a backup or restore request, you can use the **kubectl get** and the **kubectl describe** commands to show information about your request.

Managing container backup and restore jobs

You can query information about backup and restore jobs, pause and resume scheduled backup jobs, and delete snapshot and copy backups that are no longer needed.

Viewing the status of backup and restore jobs

After you submit a backup or restore request, you can use the **kubectl get** and the **kubectl describe** commands to show information about your request.

Procedure

1. To show a listing of all Kubernetes Backup Support requests in a namespace, issue the **kubectl get** command as follows:

```
kubectl get baasreq -n namespace
```

For example, to show all requests in the `production-01` namespace, issue the following command:

```
kubectl get baasreq -n production-01
```

The output is similar to the following example:

NAME	AGE
vol08-adhoc	17d
inv-adhoc2	17d
db-vol08	18d
db-vol09	17d

2. Using the results from Step [“1” on page 329](#), issue the **kubectl describe** command as follows:
 - To show the status of a job, specify the unique request name in the following command:

```
kubectl describe baasreq request_name -n namespace
```

For example, to show the status of the on-demand snapshot request named vol08-adhoc in the production-01 namespace, issue the following command:

```
kubectl describe baasreq vol08-adhoc -n production-01
```

The output is similar to the following example:

```
Name:          vol08-adhoc
Namespace:     production-01
Labels:        <none>
Annotations:   kubectl.kubernetes.io/last-applied-configuration:
               {"apiVersion":"baas.io/v1alpha1","kind":"BaaSReq","metadata":{"annotations":
               {},"name":"vol08-
               adhoc","namespace":"production-01"},"spec":{"pvcname"...
API Version:   baas.io/v1alpha1
Backupstatus:  Ready
Kind:          BaaSReq
Metadata:
  Creation Timestamp:  2019-06-11T18:59:46Z
  Generation:          5
  Resource Version:    5105243
  Self Link:           /apis/baas.io/v1alpha1/namespaces/baas/baasreqs/vol08-adhoc
  UID:                 144a1390-8c7b-11e9-8d52-005056bd89a3
Spec:
  Inprogress:  None
  Instanceid:  bmjqe7tec8ma43hm9cv0
  Origreqtype: ondemandbackup
  Pvcname:     db-vol08
  Requesttype: OnDemandBackup
  Spppvcname:  production-01:db-vol08
Status:
  Timestamp:  2019-06-11 18:59:46
  Type:       FAST-ONDEMAND
  Volumename: db-vol08
  Events:     <none>
```

- To show the list of all backups for a persistent volume claim (PVC), specify the name of the PVC and the namespace in the following command:

```
kubectl describe baasreq pvc_name -n namespace
```

For example, to show all backups for PVC db-vol08 in the production-01 namespace, issue the following command:

```
kubectl describe baasreq db-vol08 -n production-01
```

The output is similar to the following example:


```

Name:          db-vol08
Namespace:     production-01
Labels:        <none>
Annotations:   kubectl.kubernetes.io/last-applied-configuration:

{"apiVersion":"baas.io/v1alpha1","kind":"BaaSReq","metadata":{"annotations":{"name":"db-
vol08","namespace":"production-01"},"spec":{"requesttyp...
API Version:   baas.io/v1alpha1
Backupstatus:  Ready
Kind:          BaaSReq
Metadata:
  Creation Timestamp:  2019-05-23T20:17:21Z
  Generation:          3286
  Resource Version:     5105242
  Self Link:            /apis/baas.io/v1alpha1/namespaces/baas/baasreqs/db-vol08
  UID:                  c55be870-7d97-11e9-8d52-005056bd89a3
Spec:
  Inprogress:          None
  Instanceid:          bmjqe7tec8ma43hm9cv0
  Origreqtype:         backup
  Requesttype:         Backup
  Spppvname:           production-01:db-vol08
  Sla:                 test
Status:
  Timestamp:           2019-05-30 13:27:21
  Type:                FAST
  Timestamp:           2019-05-30 13:32:21
  Type:                COPY
  Timestamp:           2019-06-11 18:59:46
  Type:                FAST-ONDEMAND
  Volumename:          db-vol08
  Events:              <none>

```

- To show information about a restore job, issue the following command:

```
kubectl describe baasreq request_name -n namespace
```

where *request_name* is the request name of the restore job and *namespace* is the namespace.

Results

In the command output, the **Backupstatus** field shows the status of a backup job. For restore jobs, the **Restorestatus** field shows the status of the job. For more information, see [“Status of backup and restore jobs”](#) on page 332.

The **instanceid** field contains a randomly generated string that uniquely identifies a volume in IBM Spectrum Protect Plus.

The **Spppvname** field shows the name of the PVC that is reported in the IBM Spectrum Protect Plus **Jobs and Operations** window. The *namespace:pvc_name* format is used for the **Spppvname**. The values for the **instanceid** and **Spppvname** fields uniquely identifies a backup in IBM Spectrum Protect Plus.

In backup requests, the **Status** section shows the list of backups that were completed. For each backup, the time stamp of the backup is listed, followed by the type of backup that was run. The types of backups are defined as follows:

FAST

Denotes the backup type for a snapshot backup that is taken during a scheduled backup operation.

COPY

Denotes the backup type for a copy backup that is stored on an IBM Spectrum Protect Plus vSnap server.

FAST-ONDEMAND

Denotes the backup type for an on-demand snapshot backup.

Status of backup and restore jobs

When you use the **kubect1 describe** command to show information about backup and restore jobs, the status of backup and restore jobs is displayed in the command output.

To display the status of a specific Kubernetes Backup Support request, enter the following command:

```
kubect1 describe baasreq request_name -n namespace
```

where *request_name* is the name of the request, and *namespace* is the namespace in which the persistent volume exists. For more information, see [“Viewing the status of backup and restore jobs” on page 329](#).

Reported backup status

The status of a backup job is shown in the Backupstatus field in the command output. The following table shows the possible statuses of a backup request:

Table 34. Status of backup jobs	
Backup status	Description
None	No backup jobs were started for this schedule.
Requested	A backup job was started for this schedule.
Ready	At least one backup job was completed for this schedule.
Paused	This schedule is paused and no backup jobs were completed.
PausedReady	This schedule is paused and one or more backup jobs were completed.
AutoPaused	This schedule was paused due to scaling the owning deployment to 0 and no backup jobs were completed.
AutoPausedReady	This schedule was paused due to scaling the owning deployment to 0 and one or more backup jobs were completed.
Destroyed	All snapshot and copy backups of a persistent volume claim were deleted.
Invalid	An issue occurred with the request. A possible explanation is listed in the Errmsg field.

Reported restore status

The status of a restore job is shown in Restorestatus field in the command output. The following table shows the possible statuses of a restore job:

Table 35. Status of restore jobs	
Restore status	Description
None	No restore jobs were requested.
Requested	A snapshot or a copy backup restore job is requested.
Restored	A snapshot or a copy backup was successfully restored.

Table 35. Status of restore jobs (continued)

Restore status	Description
Invalid	An issue occurred with the request. A possible explanation is listed in the Errmsg field.

Pausing scheduled backups

You can pause a backup schedule for a container when you do not want scheduled backup jobs to run. For example, you might want to pause a backup schedule to perform maintenance tasks on a container or to restart the container.

Before you begin

Do not pause a backup schedule for a longer duration than the longest retention period in the SLA. Existing backups expire according to the SLA and you will not be able to restore the backups after they expire.

For example, assume that the SLA that is associated with a volume is **weekly**, with a retention period of 24 days for snapshots and 31 days for copy backups. You pause the schedule for 35 days and resume the schedule on day 36. No snapshots or copy backups will be available for restore jobs between days 32 and 35.

Procedure

1. Create a YAML file that contains the following properties:

```
#-----
# Filename: filename.yaml
#-----

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: pvc_name
  namespace: namespace
spec:
  requesttype: Pause
```

where:

pvc_name

The name of the backup request, which must match the name of the PVC for the volume that is being backed up.

namespace

The namespace in which the PVC exists.

2. Submit the pause request by issuing the following command:

```
kubectl apply -f filename.yaml
```

After you submit the pause request, no further backup jobs are run for the PVC until you resume the schedule. Any existing snapshots and copy backups will expire according to the SLA that is associated with the PVC.

Related tasks

[“Pausing scheduled backups” on page 333](#)

You can pause a backup schedule for a container when you do not want scheduled backup jobs to run. For example, you might want to pause a backup schedule to perform maintenance tasks on a container or to restart the container.

Related information

[“Types of requests in Kubernetes Backup Support” on page 310](#)

Resuming scheduled backups

When you want a paused backup schedule to run again, you can resume the schedule. For example, you can resume a backup schedule when a maintenance task on a container is completed and when the container is restarted.

Procedure

1. Create a YAML file that contains the following properties.

```
#-----  
# Filename: filename.yaml  
#-----  
  
apiVersion: "baas.io/v1alpha1"  
kind: BaaSReq  
  
metadata:  
  name: pvc_name  
  namespace: namespace  
spec:  
  requesttype: Resume
```

where:

pvc_name

The name of the backup request, which must match the name of the PVC for the volume that is being backed up.

namespace

The namespace in which the PVC exists.

2. Submit the resume request by issuing the following command:

```
kubectl apply -f filename.yaml
```

After you submit the request, the scheduled backup operations will resume within the window that is defined by the SLA policy.

Related tasks

[“Pausing scheduled backups” on page 333](#)

You can pause a backup schedule for a container when you do not want scheduled backup jobs to run. For example, you might want to pause a backup schedule to perform maintenance tasks on a container or to restart the container.

Related information

[“Types of requests in Kubernetes Backup Support” on page 310](#)

Deleting container backups

You can delete snapshot and copy backups of a persistent volume claim (PVC) by submitting a **destroy** request.

Before you begin

Before you submit a **destroy** request to delete container backups, consider the following consequences:

- All snapshots of the PVC will be deleted in the Kubernetes environment. These snapshots include snapshots that were created by on-demand and scheduled backups, and snapshots that were manually created.
- The copy backups on the IBM Spectrum Protect Plus vSnap server will be marked for deletion. The deletion is managed by IBM Spectrum Protect Plus.
- The original backup request will not be deleted by the **destroy** request. You must run the **kubectl delete** command to delete it.

- The **destroy** request is not supported for on-demand backups. Use the **kubect1 delete** command to delete an on-demand backup request. An on-demand snapshot is deleted when the snapshot expires or when the scheduled backup is destroyed.

Procedure

1. Create a YAML file for the **destroy** request that contains the following properties:

```
#-----  
# Filename: filename.yaml  
#-----  
  
apiVersion: "baas.io/v1alpha1"  
kind: BaaSReq  
  
metadata:  
  name: pvc_name  
  namespace: namespace  
spec:  
  requesttype: Destroy
```

where:

filename

The name of the YAML configuration file.

pvc_name

The name of the request, which must match the name of the PVC that was backed up. For example, if you want to delete all snapshots and copy backups for the PVC named `db-vo101`, the name of the request must also be `db-vo101`.

namespace

The namespace in which the PVC exists.

2. Submit the **destroy** request by entering the following command on the command line:

```
kubect1 apply -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

3. To check that the snapshots and copy backups for a PVC are deleted, issue the following command:

```
kubect1 describe baasreq pvc_name -n namespace | grep Backupstatus
```

In the command output, the following status shows that the backups were deleted:

```
Backupstatus: Destroyed
```

What to do next

As a best practice, delete the completed request by issuing the following command:

```
kubect1 delete baasreq pvc_name -n namespace
```

Deleting completed requests has the following benefits:

- It reduces the size of the etcd database and allows you to reuse the name of a request for another operation.
- It makes troubleshooting easier.
- It makes it easier for you to track backup and restore requests that are running in your Kubernetes cluster.
- At any point in time, you have a clear picture of requests that are running in on your cluster when you issue the following command:

```
kubect1 get baasreq -n namespace
```

If you delete the backup request without first destroying the backup, the backup request will continue to run and backups will be made according to specified SLA policy until Kubernetes Backup Support is restarted or the `baas-etcd-client` pod is restarted.

If you accidentally deleted the backup request without destroying the backup first, you must complete the following actions:

1. Manually restart the `baas-etcd-client` pod.
2. Manually delete the volume snapshots for the volume.

Related information

[“Types of requests in Kubernetes Backup Support” on page 310](#)

Monitoring Kubernetes Backup Support jobs and running reports

The backup administrator can use the IBM Spectrum Protect Plus user interface to monitor Kubernetes Backup Support jobs and create reports that show the backup history of containers.

Viewing job logs

You can use the **Jobs and Operations** window to monitor Kubernetes Backup Support jobs, review job history, and view scheduled jobs.

About this task

Container-specific jobs can be identified by the `kubernetesvol` prefix, the name of the persistent volume claim (PVC), or the CTGGK message identifier.

For example, the name of a job can be in the following format:

```
kubernetesvol_internalID_namespace:pvcname
```

where:

internalID

A random string generated by Kubernetes Backup Support that uniquely identifies a volume.

namespace

The namespace in which a PVC exists.

pvcname

The name of PVC that is being protected.

Procedure

1. In the IBM Spectrum Protect Plus navigation pane, click **Jobs and Operations**.
2. Click the appropriate tab:
 - To show the backup and restore jobs that are running, click **Running Jobs**.
 - To show the jobs that ran successfully, completed processing with warnings, or jobs that failed, click **Job History**. You can download a job log from the page by selecting the job and clicking **Download.zip**.
 - To view the status of scheduled jobs, click **Schedule**.

Related concepts

[“Managing jobs and operations” on page 343](#)

You can manage and monitor jobs in the **Jobs and Operations** window. You can also configure scripts to run before or after jobs.

Creating backup history reports

You can run a report to show the backup history of your protected persistent volumes. By viewing the backup history, you can determine the health of your backups.

About this task

For each persistent volume claim (PVC), the backup history shows information about the Container Storage Interface (CSI) snapshots that were created in the Kubernetes environment and the backups that were copied to the IBM Spectrum Protect Plus vSnap server. You can view information such as the date and time of the backup operation, the size of the backup, and how fast it took to process the copy backups. Because all snapshots and copy backups are recorded in IBM Spectrum Protect Plus, you can see whether your scheduled backups are running according to the service level agreement (SLA) policy that you set for the PVC.

Procedure

1. In the IBM Spectrum Protect Plus navigation pane, click **Reports and Logs**.
2. Expand the **Protection** list and take one of the following actions:
 - To create a new report, click **Container Persistent Volume Backup History**.
 - To run a report that you saved previously, expand the **Container Persistent Volume Backup History** list and click the name of a saved report.
3. In the **Options** section, take one of the following actions:
 - To run the report immediately with the default parameters or the parameters of a saved report, click **Run**.
 - To customize a new report, update the parameters in the **Options** section and click **Run**. You can also provide a name and description and save the definitions for future use.
 - To change the definitions of a saved report, update the parameters in the **Options** section and click **Save**. Then, click **Run**.

The backup history report is displayed in the **Container Persistent Volume Backup History** section of the window. To download the report, click **Download**.
4. Optional: To schedule the report and send it to a recipient, select **Define Schedule**.
 - a) In the **Frequency** field, specify how often to run the report.
 - b) In the **Start Time** field, specify the date and time for when to start running the report.
 - c) In the email recipient field, enter at least one email address and click **Add a recipient**. The email address must be a valid address.
 - d) Click **Save**.

Results

The backup history report is shown in the **Container Persistent Volume Backup History** section of the window. The descriptions of the reported data are shown in the following table:

Table 36. Details of the backup history report

Column	Description
SLA Policy	<p>The SLA policy that is used to protect a PVC.</p> <p>Because only one SLA can be assigned to a volume, the SLA can contain the following elements:</p> <pre>instancename_internalID_namespace:pvcname</pre> <p>The <i>instancename</i> value is composed of the namespace and the PVC name.</p> <p>The <i>internalID</i> value is a random string generated by Kubernetes Backup Support that uniquely identifies a volume.</p>
Protection Time	The date and time when each backup job was completed.
Status	The status of each backup. If a backup job failed, a possible reason is provided.
Snapshot Backup?	An indication of whether the backup instance is a snapshot backup. A check mark is displayed in the column to indicate that the instance is a snapshot backup. When a check mark is displayed, no data is shown in the Backup Size and Backup Speed columns.
Backup Size	For copy backups, the amount of data that was backed up to the vSnap server. For snapshot backups that were created in the Kubernetes environment or for backups that failed, no size is shown.
Backup Speed	The rate at which a copy backup was completed. For snapshot backups or backups that failed, no data is shown.

Related concepts

[“Managing reports and logs” on page 353](#)

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 12. Protecting IBM Spectrum Protect Plus

Protect the IBM Spectrum Protect Plus application by backing up the underlying databases for disaster recovery scenarios. Configuration settings, registered resources, restore points, backup storage settings, and job information are backed up to a vSnap server that is defined in the associated SLA policy.

Backing up the IBM Spectrum Protect Plus application

Back up IBM Spectrum Protect Plus configuration settings, SLA policies, registered resources, backup storage settings, restore points, and imported keys and certificates to a vSnap server that is defined in the associated SLA policy.

Before you begin

Ensure that an appropriate SLA policy is available. To optimize backup jobs, create SLA policies 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” on page 145](#).

Restriction: You cannot select the onboard vSnap server as the target of the SLA policy for backing up IBM Spectrum Protect Plus. The onboard vSnap server is named localhost and is automatically installed when the IBM Spectrum Protect Plus appliance is initially deployed. Select a secondary external vSnap server as the target when you create the SLA policy to back up IBM Spectrum Protect Plus.

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 IBM Spectrum Protect Plus data:

1. In the navigation pane, 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 99](#).

Restoring the IBM Spectrum Protect Plus application

Restore IBM Spectrum Protect Plus configuration settings, restore points, and job information that were backed up to the vSnap server. The data can be restored to the same location or another IBM Spectrum Protect Plus location.

About this task



Attention: An IBM Spectrum Protect Plus restore operation overwrites all data in the IBM Spectrum Protect Plus virtual appliance or alternate virtual appliance location. 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.

If restoring a cloud backup, the cloud resource or repository server must be registered on the alternate IBM Spectrum Protect Plus location.

When a catalog restore job is started, a job session identifier (ID) is assigned. During the initial phase, the job will be available to be monitored in the IBM Spectrum Protect Plus UI on the job management screen until the recovery step initiates the internal database restore. Once 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/adminconsole/managedb-catalogrestore-time.log`, where *time* is epoch time. Data contained in this log is relates to the restore of the mongo configuration and recovery catalog. After the process is complete, the *virgo* service will start and the data is written to the *virgo* log. When the job is complete, the IBM Spectrum Protect Plus user interface is again accessible.

Procedure

To restore IBM Spectrum Protect Plus data:

1. In the navigation pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore**.
2. Select a vSnap server, cloud resource, or repository server.

Data can be restored to the same location, or an alternate location in disaster recovery scenarios.

Available snapshots for the server are displayed.

3. Click **Restore** for the catalog snapshot that you want to restore.
4. 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.

5. Click **Restore**.
6. To run the restore job, in the dialog box, click **Yes**.

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 343](#)

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 pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore Point Retention**.

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 391](#).

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.
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 pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore Point Retention**.
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 391](#).

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 343](#)

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Chapter 13. 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.

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:

Related concepts

[“Protecting hypervisors” on page 151](#)

You must register the hypervisors that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the virtual machines and resources that are associated with the hypervisors.

[“Protecting applications” on page 189](#)

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” on page 145](#)

You can create custom SLA policies to define backup frequency, retention, replication, and copy policies that are specific for your environment.

[“Running an ad hoc backup job” on page 349](#)

With an ad hoc backup job, you can back up one or more resources that are associated with an SLA policy on demand.

Concurrent jobs

Jobs that overlap other jobs are referred to as concurrent jobs

To determine whether a job is running or ran concurrently with another job, click **Jobs and Operations > Job History**, select a job, and click **Concurrent Jobs**.

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.

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 Chapter 9, “Protecting hypervisors,” on page 151 and Chapter 10, “Protecting applications,” on page 189 . For example, the backup topic for VMware is “Backing up VMware data” on page 155 .

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 349.
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 Chapter 9, “Protecting hypervisors,” on page 151 and Chapter 10, “Protecting applications,” on page 189. For example, the restore topic for VMware is “Restoring VMware data” on page 164.

Related concepts

[“Job types”](#) on page 343

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Related tasks

[“Creating an SLA policy”](#) on page 145

You can create custom SLA policies to define backup frequency, retention, replication, and copy policies that are specific for your environment.

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 pane, click **Jobs and Operations**, and click the **Schedule** tab.
2. Choose the job that you want to run and click **Actions > 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 and job logs

You can view information about the status of your running jobs, and the overall status of the jobs that are completed. Assess the completion of jobs that completed with failures or warnings, view the associated job logs, and rerun the jobs.

About this task

Procedure

To view jobs and job logs, complete the following steps:

1. In the navigation pane, click **Jobs and Operations**.
2. In the **Running Jobs** page, view the status of the jobs that are currently running, as shown in the following picture.

Jobs and Operations Create job

Running Jobs Job History Active Resources Schedule

8 Total Jobs 0 Backup 0 Inventory 0 Maintenance 8 Restore

CPU Usage 29% IBM Spectrum Protect Plus Host Machine

Sort By Start ↑↓ Search by name.. Q ▼

onDemandRestore_1571250656746
Exchange
Type: Restore | Activity: Resource active
Start Time: Oct 16, 2019 7:30:57 PM
Duration: 0h 16m 7s Databases Completed: 0/1

ce9_db2_recurring
Type: Restore | Activity: Resource active
Start Time: Sep 11, 2019 9:00:00 PM
Duration: 840h 0m 33s Databases Completed: 1/1

onDemandRestore_1567562307977
SQL

onDemandRestore
Type: Restore | Start Time: Oct 16, 2019 7:30:57 PM

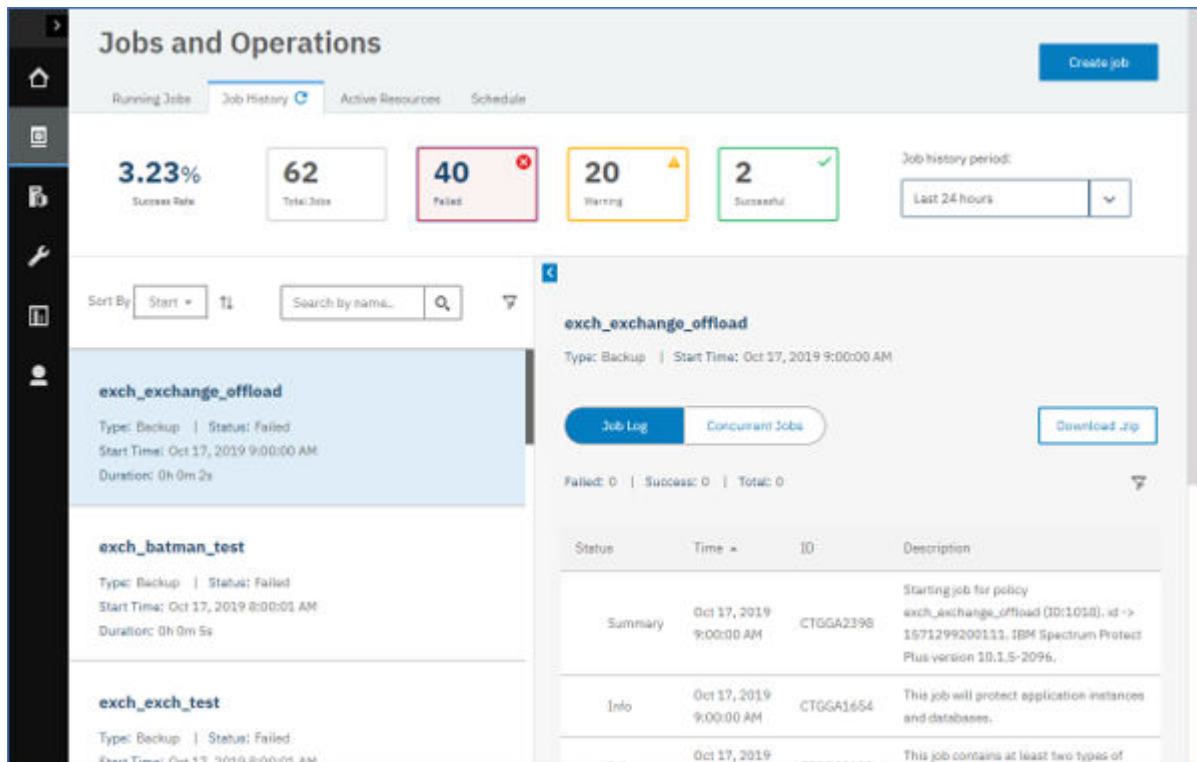
Job Log Concurrent Jobs Download .zip



Failed: 0 | Success: 0 | Total: 1

Status	Time	ID	Description
Summary	Oct 16, 2019 7:30:57 PM	CTGGA2398	Starting job for policy onDemandRestore_1571250656746 (ID:1274). id -> 1571250657082. IBM Spectrum Protect Plus version 10.1.5-2096.
Detail	Oct 16, 2019 7:30:57 PM	CTGGA2109	Policy has (1) destination database mappings.

3. To view completed jobs, click **Job History**.

The ribbon across this screen shows the status of historical jobs. Use the filter to define the duration of the job history to display. Clicking the job displays the job log for that job on the screen, and you can choose to download the log for more details.



4. To view the active resources in your environment, click **Active Resources**.
Shows application and hypervisor active resources.
5. To view the overall schedule for all jobs, click **Schedule**.
Using the **Actions** menu, you can choose to start a job or pause a schedule. You can also edit some recurring and maintenance job schedules by clicking the schedule icon, , and saving your changes.
To edit a restore job, click the edit icon for that job, .

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 pane, click **Jobs and Operations**, and click the **Schedule** tab.
2. Choose the job that you want to pause, and click **Actions > Pause Schedule**.
3. To resume the job schedule, click **Actions > 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 "Snapshot 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 pane, 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 "Snapshot 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 hypervisors," on page 151 and Chapter 10, "Protecting applications," on page 189.
	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 pane, 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 pane, 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 pane, 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 back up one or more resources that are associated with an SLA policy on demand.

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.


Procedure


To run an ad hoc backup job, complete the following steps:

1. In the navigation pane, 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 pane 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, select the SLA policy and then 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 pane 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 pane, 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 351 .

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 pane, 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 14. 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.

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 pane, click **Reports and Logs > Reports**.
2. Expand **Backup Storage Utilization** in the **Reports** pane.

The following reports are available:

VM Backup Utilization

Review the utilization of your virtual machine (VM) backups on backup storage, including the following data:

- The name of each VM, its location, and the associated hypervisor.
- The SLA policy that is used to protect the VM.
- The location of the backup storage. The backup storage can be the host name or IP address of a disk, the name of a cloud server, or the name of the repository server.
- The size of each VM backup.
- The number of restore points that are available for each VM.

For VMware virtual machines, to narrow your results to show VMs that have VMware tags, select one or more available tags in the **Tags** drop-down menu. The default value is **All**, which shows data for all VM backups.

vSnap Storage Utilization Report

Review the storage utilization of your vSnap servers, including the availability status, free space, and used space. 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.

Use the report options to filter specific vSnap servers to display. 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/deduplication ratio of the vSnap server.

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 359](#)

You can run, save, or schedule reports in IBM Spectrum Protect Plus.

[“Types of reports” on page 353](#)

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 pane, click **Reports and Logs > Reports**.
2. Expand **Protection** in the **Reports** pane.

The following reports are available:

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 hypervisors 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**.

The **Summary View** displays an overview of your virtual machine protection status, including the number of unprotected and protected virtual machines and the managed capacity of the protected virtual machines. The managed capacity is the used capacity of a virtual machine. The **Detail View** provides further information about the protected and unprotected virtual machines, including names and location.

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** provides further information about the protected and unprotected databases, including their names and location.


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 by failed or successful jobs and time of the last backup. The report can be further filtered by specific service level agreement (SLA) policies. In the **Detail View**, click the plus icon  next to an associated job to view job details, such as the reason why a job failed or the size of a successful backup.

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 by failed or successful jobs and time of the last backup. The report can be further filtered by specific SLA policies. In **Detail View**, click the plus icon  next to an associated job to view further job details, such as the reason why a job failed or the size of a successful backup.

VM SLA Policy RPO Compliance report

The VM SLA Policy RPO Compliance report displays virtual machines in relation to recovery point objectives as defined in SLA policies. The report displays the following information:

- Virtual machines in compliance
- Virtual machines not in compliance
- Virtual machines in which the last backup job session failed

Use the report options to filter by hypervisor type and to select specific hypervisors to display. The report can be further filtered by virtual machines that are in compliance or not in compliance with the defined RPO.

Database SLA Policy RPO Compliance report

The Database SLA Policy RPO Compliance report displays databases in relation to recovery point objectives as defined in SLA policies. The report displays the following information:

- Databases in compliance
- Databases not in compliance
- Databases in which the last backup job session failed

Use the report options to filter by application type and to select specific application servers to display. The report can be further filtered by databases that are in compliance or not in compliance with the defined RPO, or by protection type, including data that was backed up to vSnap or by using replication.

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 and by specific service level agreement (SLA) policies. Set a number in the **Backup History for Past Number of Days** field to show the backup history for a specified number of days. The default value is 1.

Tip: Optionally, the custom report can be saved and run at a future time by giving the report a name and description in the **Save Custom Report** pane.

Related concepts

[“Types of reports” on page 353](#)

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 pane, click **Reports and Logs > Reports**.
2. Expand **System** in the **Reports** pane.

The following reports are available:

Configuration report

Review the configuration of the application servers, hypervisors, and backup storage that is available. Use the report options to filter the configuration types to display. The report displays the name of the resource, resource type, associated site, and the SSL connection status.

Job report

Review the available jobs in your configuration. Run this report to view jobs by type, their average duration, and their successful run percentage. Use the report options to filter the job types to display and to display jobs that ran successfully over a period of time. The **Summary View** lists jobs by type along with the number of times a job session is run, completed, or failed. Job sessions listed as Other 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 virtual machines that are protected by a backup job, the average run time, and the next scheduled run time if the job is scheduled.

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 353](#)

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 pane, click **Reports and Logs > Reports**.
2. Expand **VM Environment** in the **Reports** pane.
3. Choose the VM report you want to run from the following list.

VM report type	Options
VM Datastores	<p>Choose this report to review the VM and storage usage of datastores in your VM environment. This report shows the datastore count, capacity, and free space for each hypervisor. To run this report, take the following actions:</p> <ol style="list-style-type: none">a. In the Hypervisor Type section, select the type of hypervisors to include in the report. The default value is All.b. In the Hypervisor section, select the hypervisors to include in the report. The default value is All.c. In the Detail View Filter section, select the threshold of percentage usage to include in the report. The default value is >80% space used. If you keep the default value, the report will include only the hypervisors in which at least 80% of space is used.d. Click Show only Orphaned Datastores to view datastores that do not have any VMs assigned to them, or to view VMs that are inaccessible.e. Proceed to Step 4.
VM Storage	<p>Choose this report to review the provisioned space in your datastores and hypervisors in your VM environment. Take the following actions:</p> <ol style="list-style-type: none">a. In the Hypervisor Type section, select the type of hypervisors to include in the report. The default value is All.b. In the Hypervisor section, select the hypervisors to include in the report. The default value is All.c. Proceed to the Step 4.
VM LUNs	<p>Choose this report to review the storage usage of your VM logical unit numbers (LUNs). The report includes details for each LUN in your environment with associated datastores per volume, capacities, transport type, and RDM. Take the following actions:</p>

VM report type	Options
	<ul style="list-style-type: none"> a. In the Hypervisor Type section, select the type of hypervisors to include in the report. The default value is All. b. In the Hypervisor section, select the hypervisors to include in the report. The default value is All. c. Click Show only Orphaned Datastores to view datastores that do not have any VMs assigned to them, or to view VMs that are inaccessible. d. Proceed to Step 4.
VM Snapshot Sprawl	<p>Choose this report to review the details of the snapshots for your VM environment. Take the following actions:</p> <ul style="list-style-type: none"> a. In the Hypervisor Type section, select the type of hypervisors to include in the report. The default value is All. b. In the Hypervisor section, select the hypervisors to include in the report. The default value is All. c. Choose a Snapshot Creation Time. The default value is >1 Year. If you keep the default value, the report shows snapshots that were created over 1 year ago. d. Select tags that are associated with a particular VM you are looking for from the Tags menu. e. Proceed to Step 4.
VM Sprawl	<p>Choose this report to review the status of your VMs. The report lists the VMs in your environment that are powered off, powered on, or suspended. It also lists VM templates. Take the following actions:</p> <ul style="list-style-type: none"> a. In the Hypervisor Type section, select the type of hypervisors to include in the report. The default value is All. b. In the Hypervisor section, select the hypervisors to include in the report. The default value is All. c. Specify dates for the report by entering the following details: <ul style="list-style-type: none"> • Days Since Last Powered Off the default is Any. • Days Since Last Suspended the default is Any. • Days Since Last Powered On the default is >180 days.

VM report type	Options
	<p>d. Select tags that are associated with a particular VM you are looking for from the Tags menu. The default value is All.</p> <p>e. Proceed to Step 4.</p>

4. Optional: To save the report, enter a **Name** and a **Description**, and click **Save** before you run the report.
5. Optional: To run the report regularly, click **Define Schedule** in the **Schedule Report** section of the window.
Specify the **Frequency** of the report, the **Start Time**, and enter email addresses for recipients of this report.
6. Click **Run** to generate the report, which is displayed in the lower part of the window.
Tip: You might have to scroll down to see the report.
7. Optional: Download the report to your computer.

Related concepts

[“Types of reports” on page 353](#)

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Report actions

You can run, save, 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 367](#).

Procedure

To run a report, complete the following steps:

1. In the navigation pane, click **Reports and Logs > Reports**.
2. Expand a report type and select a report to run.
3. Run the report either with custom parameters or default parameters:
 - To run the report with custom parameters, set the parameters in the **Options** section, 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 353](#)

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 pane, click **Reports and Logs > Reports**.
2. Select a predefined report.
3. Set your customized parameters.
4. Define the report to run in one of the following circumstances:
 - Run on demand.
 - Create a schedule to run the report as defined by the parameters of the schedule.
5. Save the report with a customized name.

What to do next

Run the report and review the report in the **Reports** pane.

Related concepts

[“Managing reports and logs” on page 353](#)

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 pane, click **Reports and Logs > Reports**.
2. Expand a report type and select a report.
3. In the **Schedule Report** section, click **Define Schedule**.
4. Define the frequency and start time for the schedule.
5. Enter an address to receive the scheduled report in the email field, and then click **Add a recipient**.
6. Click **Save**.

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 353](#)


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 pane, 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 371](#)

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.

Enabling the display of alerts in the system log

Alerts that are generated by IBM Spectrum Protect Plus can be logged to the system log. After you enable this feature, you can search the system log to find alerts.

About this task

An alert list in the IBM Spectrum Protect Plus user interface displays errors and warnings, which can help you to identify and resolve issues. Each item in the alert list includes a status icon, a summary of the issue, and a link to associated logs. These alerts can also be logged to the system log on the IBM Spectrum Protect Plus appliance.

Procedure

To enable logging of IBM Spectrum Protect Plus alerts to the system log, complete the following steps:

1. Ensure that there are no running jobs on the IBM Spectrum Protect Plus appliance. Information for running jobs can be found on the **Jobs and Operations > Running Jobs** pane.
2. Stop the Virgo service by issuing the following command:

```
sudo systemctl stop virgo
```

3. Log in to the IBM Spectrum Protect Plus appliance by using the Secure Shell (SSH) protocol.
4. Open the persistence properties file in the following directory:

```
/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.persistence.properties
```

Tip: You can edit the persistence properties file by using the `vi` text editor, which is packaged along with the IBM Spectrum Protect Plus appliance.

5. Set the `alertInSystemLogMessages` variable to **yes**:

```
mongodb.alertInSystemLogMessages=yes
```

6. Save the properties file.

7. Restart the Virgo service by issuing the following command:

```
sudo systemctl start virgo
```

Results

Alerts are logged to the system log.

Example

```
[root@veguardian-cb3 log]# grep SPPAlert /var/log/messages
Jan 25 23:53:03 veguardian-cb3 virgo[993]: ERROR [SPPAlert] vmware_Copy_Protect:1548460344227
Job vmware_Copy_Protect (id=1006, session=1,548,460,344,227) failed.
```

Note: The following format is used for alerts that are logged by IBM Spectrum Protect Plus:

Message_type: The type can be ERROR, WARNING, and other types.

Message_source: Indicates the source of the message. This appears as SPPAlert for IBM Spectrum Protect Plus messages.

Message_text: Contains the IBM Spectrum Protect Plus alert details.

What to do next

Optionally, to search for an alert, log in to the IBM Spectrum Protect Plus appliance and search for SPPAlert.

Chapter 15. 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.

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.

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 pane, 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 365.</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 368 .

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:

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.
Application	<ul style="list-style-type: none"> • Db2 • 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.
Application Server	<ul style="list-style-type: none"> • Db2 • SQL • Oracle 	Used to grant access to application servers in IBM Spectrum Protect Plus without access to individual databases.
Hypervisor	<ul style="list-style-type: none"> • VMware • Hyper-V 	Used to grant access to hypervisor resources.
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.

Resource Type	Subtype	Description
SLA Policy	None	Used to grant access to SLA Policies for Backup operations.
System	Identity	Used to grant access to the credentials required to access your resources. Identity functionality is available through the System > Identity pane.
System Configuration	Disk	Used to grant access to vSnap backup storage servers.
System Configuration	LDAP	Used to grant access to LDAP servers for user registration.
System Configuration	Logs	Used to grant access to viewing and downloading Audit and System logs.
System Configuration	Script	Used to grant access to uploaded prescripts and postscripts.
System Configuration	Script Server	Used to grant access to script servers, where scripts are run during a Backup or Restore job.
System Configuration	Site	Used to grant access to sites, which are assigned to vSnap backup storage servers.
System Configuration	SMTP	Used to grant access to SMTP servers for job notifications.
System Configuration	VADP Proxy	Used to grant access to VADP proxy servers.

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 pane, 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 pane, 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 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.

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.

The following user account roles are available:

Application Admin

The Application Admin role allows users to 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

The Backup Only role allows users to 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**.

Restore Only

The Restore Only role allows users to 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

The Self Service role allows users to 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 admin user:

- Modify and delete user accounts
- Change user passwords
- Assign user roles

An administrator can also access the administrative console by selecting **IBM Spectrum Protect Plus** from the **Authentication Type** list in the console login window and entering administrator credentials.

From the administrative console, the administrator can apply software updates, restart the IBM Spectrum Protect Plus appliance, and set the local time zone.

For more information about using the Administrative Console, see [“Logging on to the administrative console” on page 135](#).

VM Admin

The VM Admin role allows a users to 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 pane, 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.

Option	Actions
	To view available permissions and their usage, see “Permission types ” on page 369.

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 SSL certificates to access cloud servers.
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.

Name	Permissions	Description
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.
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.
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 pane, 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 pane, 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 364 and [“Creating a role”](#) on page 368.

Procedure

To create an account for an individual user, complete the following steps:

1. In the navigation pane, 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 name and password for the user.
5. In the **Assign Role** section, select one or more roles for the user.
6. In the **Permission Groups** section, review the permissions and resources that are available to the user, and then click **Continue**.
7. In the **Add Users - Assign Resources** section, assign one or more resource groups to the user, and then click **Add resources**.
The resource groups are added to the **Selected Resources** section.
8. 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 127](#).
- 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 368](#) and [“Creating a resource group” on page 364](#).

Procedure

To create a user account for an LDAP group, complete the following steps:

1. In the navigation pane, 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. In the **Assign Role** section, select one or more roles for the user.
7. In the **Permission Groups** section, review the permissions and resources that are available to the user, and then click **Continue**.
8. In the **Add Users - Assign Resources** section, assign one or more resource groups to the user, and then click **Add resources**.
The resource groups are added to the **Selected Resources** section.
9. 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 users who are assigned to the SUPERUSER role. If a user is a member of the SUPERUSER role, you can change only the password for the user.

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 pane, 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 settings

Edit the user name and password, associated roles, and resource groups.

Modify resources

Edit the associated resource groups.


4. Modify the settings for the user, and then click **Update user** or **Assign resources**.

Deleting a user account

You can delete any user account, with the exception of users who are assigned to the SUPERUSER role.

Procedure

To delete a user account, complete the following steps:

1. In the navigation pane, click **Accounts > User**.
2. Select a user.
3. Click the options icon , and then click **Delete user**.

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 pane, 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 pane, 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 pane, click **Accounts > Identity**.
2. Click the delete icon  that is associated with an identity.
3. Click **Yes** to delete the identity.

Chapter 16. Licensing

License auditing in IBM Spectrum Protect Plus is enabled by default to determine if the current usage is within license entitlement levels and to prevent potential license violations.

IBM Spectrum Protect Plus generates entitlement audit logs as IBM® Software License Metric Tag (.slmtag) files. IBM® License Metric Tool (ILMT) is then used to translate the file and generate License Consumption Reports. Use the information in this section to interpret your .slmtag files.

Software License Metric (SLM) tags

IBM Spectrum Protect Plus generates entitlement audit logs as IBM® Software License Metric Tag (.slmtag) files. IBM® License Metric Tool (ILMT) is then used to translate the file and generate License Consumption Reports. Use the provided information to interpret your .slmtag files.

The .slmtag files can store information up to a maximum file size of 1 MB, after which the file is archived and a new log file is created. A maximum of 10 log files are kept.

Upgrade requirements: If you are upgrading IBM Spectrum Protect Plus from a prior release, you must run the maintenance job to update existing .slmtag files.

Log format

The .slmtag files are stored in XML format, with new metric records appended to the end of the file.

The following is a sample .slmtag file:

```
<SchemaVersion>2.1.1</SchemaVersion>
<SoftwareIdentity>
  <SoftwareIdentity name>"IBM Spectrum Protect Plus"</Name>
  <InstanceId>/opt/virgo</InstanceId>
</SoftwareIdentity>
<Metric logTime="2018-11-05T16:05:09+00:00">
  <Type>HYPERVISOR_SERVER_COUNT</Type>
  <SubType>HYPERVISOR_SERVER_COUNT</SubType>
  <Value>0</Value>
  <Period>
    <StartTime>2018-11-05T16:05:09+00:00</StartTime>
    <EndTime>2018-11-05T16:05:09+00:00</EndTime>
  </Period>
</Metric>
<Metric logTime="2018-11-05T16:05:09+00:00">
  <Type>APPLICATION_INSTANCE_COUNT</Type>
  <SubType>APPLICATION_INSTANCE_COUNT</SubType>
  <Value>0</Value>
  <Period>
    <StartTime>2018-11-05T16:05:09+00:00</StartTime>
    <EndTime>2018-11-05T16:05:09+00:00</EndTime>
  </Period>
</Metric>
```

where the Value element displays the number of hosts in all the resource groups with packages deployed for an instance group, at the specified time in the EndTime element.

The file grows over time and may be edited to remove older metric elements. Ensure that you retain elements long enough for ILMT scanning; the scanning frequency is determined by the ILMT administrator, but generally it should be sufficient to keep elements for a month.

Log location

The .slmtag file is located in the /data/slmtag directory.

Related concepts

[“Job types” on page 343](#)

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Related tasks

[“Starting jobs on demand” on page 345](#)

You can run any job on demand, even if the job is set to run on a schedule.

Integration with IBM License Metric Tool (ILMT)

Use IBM License Metric Tool (ILMT) to help determine whether your system environment is compliant with licensing requirements.

ILMT provides useful features for managing virtualized environments and measuring license utilization. ILMT discovers the software that is installed in your infrastructure, helps you to analyze the consumption data, and allows you to generate audit reports. Each report provides you with different information about your infrastructure, for example the computer groups, software installations, and the content of your software catalog.

By default, every ILMT audit report presents data from the previous 90 days. You can customize the type and amount of information displayed in a report by using filters, and save your personal settings for future use. You can also export the reports to .csv or .pdf format, and schedule report emails so that specified recipients are notified when important events occur.

For more information, see the [IBM License Metric Tool](#) product documentation.

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 2014120](#).

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.

For information about log files, see the protecting applications or protecting hypervisors backup sections.

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.

Troubleshooting Kubernetes Backup Support

To help troubleshoot issues with Kubernetes Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Troubleshooting quick reference

Solutions to basic Kubernetes Backup Support problems are provided.

Use the solutions in the following table to resolve basic problems that might occur with Kubernetes Backup Support operations. If you still cannot resolve a problem, see [“Troubleshooting Kubernetes Backup Support operations”](#) on page 379 for more detailed troubleshooting procedures.

Table 37. Solutions to basic problems

Problem	Solution
<p>The Kubernetes Backup Support request is invalid. For example, the Backupstatus or Restorestatus field is listed as Invalid when you run the following command:</p> <pre>kubect1 describe baasreq request_name -n namespace</pre> <p>where:</p> <p>request_name The name of the backup or restore request. For backup requests, the value is the name of the persistent volume claim (PVC). For restore requests, the name must be unique, and must not be the same as the name of the PVC.</p> <p>namespace The namespace in which the PVC exists.</p>	<p>Ensure that the request is structured correctly by verifying the following elements in the YAML file:</p> <ul style="list-style-type: none"> • Ensure that there are no typographical errors. • Ensure that the correct case is used in the statements. Kubernetes is case sensitive. For example, ensure that the API version declaration is listed as <code>apiVersion</code> and not <code>apiversion</code>. • For restore requests: <ul style="list-style-type: none"> – Ensure that the time stamp for a restore point is specified correctly in the restorepoint field. – Ensure that the restore type is specified correctly in the restoretype field. <p>For more information, see “Restoring container data” on page 326.</p>
<p>The snapshots are failing.</p>	<p>Take one or more of the following actions:</p> <ul style="list-style-type: none"> • Verify the Ceph-CSI configuration to ensure that your containers are running correctly. The CSI software is required for snapshot backups. • Ensure that a volume snapshot class is defined for the PVCs that are being backed up. • Ensure that the secret is in the correct namespace (the namespace for the PVC). • Ensure that the configurations are correct in the ConfigMap (baas-configmap). <p>For more information, see “Troubleshooting issues with snapshot backup jobs” on page 380.</p>
<p>The data mover fails to start.</p>	<p>Take one or more of the following actions:</p> <ul style="list-style-type: none"> • Ensure that the Ceph RBD volume is mounted. You can verify whether the Ceph RBD volume is failing to mount by issuing the kubect1 describe command on the data mover pod. • In the output of the kubect1 describe command, check the events to ensure that the volume has been initialized by running the PVC as part of another pod in read/write mode. • In the output of the kubect1 describe command, check for authentication failure events. To resolve authentication errors, ensure that you are running a secure Docker registry. Ensure that the pull secret is in the namespace of the PVC. For instructions, see Pull an Image from a Private Registry.

Table 37. Solutions to basic problems (continued)

Problem	Solution
Access is denied or the connection fails while mounting NFS volumes from the vSnap server.	<p>Take one or more of the following actions:</p> <ul style="list-style-type: none"> • Check the data mover network policy. Ensure that the vSnap server addresses match the IBM Spectrum Protect Plus server addresses. • Ensure that a direct connection from the Kubernetes cluster to the IBM Spectrum Protect Plus vSnap server exists. Connection by proxies is not supported.
The scheduler, transaction manager, and controller pods have started but each pod continues to restart. In the output of the kubectl describe command for the transaction manager pod, the events indicate that the liveness probe failed.	<p>Verify that the values for the CLUSTER_API_SERVER_IP_ADDRESS and CLUSTER_API_SERVER_PORT parameters are correctly specified in the baas_config.cfg configuration file.</p> <p>If you update the values in the baas_config.cfg file, issue the following command to update the configuration:</p> <pre>./baas_install.sh -u</pre> <p>Alternatively, you can uninstall and reinstall Kubernetes Backup Support to clear the previous log files. For instructions, see “Uninstalling Kubernetes Backup Support” on page 320 and “Installing and deploying Kubernetes Backup Support images” on page 315.</p>

Related tasks

[“Collecting Kubernetes Backup Support log files for troubleshooting”](#) on page 385

You can generate debugging log files in the Kubernetes environment to troubleshoot the deployment of Kubernetes Backup Support and Kubernetes Backup Support operations on the IBM Spectrum Protect Plus server.

Troubleshooting Kubernetes Backup Support operations

Troubleshooting procedures are available to help you diagnose and resolve Kubernetes Backup Support issues.

The following instructions are provided:

- [“Viewing log files”](#) on page 379
- [“Troubleshooting issues with snapshot backup jobs”](#) on page 380
- [“Troubleshooting issues with copy backup jobs”](#) on page 381
- [“Troubleshooting restore jobs”](#) on page 383

Viewing log files

To troubleshoot Kubernetes Backup Support issues, start by viewing information in the log files. Log files are available for the transaction manager, controller, and scheduler components of Kubernetes Backup Support.

To view the log file for the transaction manager component, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-transaction-manager/ {print $1;exit}') -n baas -c baas-transaction-manager -f
```

To view the log file for the transaction manager worker, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-transaction-manager/ {print $1;exit}') -n baas -c baas-transaction-manager-worker -f
```

To view the log file for the controller component, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-controller/ {print $1;exit}') -n baas -f
```

To view the log file for the scheduler component, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-scheduler/ {print $1;exit}') -n baas -f
```

Tip: To help speed up the display of log files, you can add the **--since=duration** flag to the **kubectl logs** command to return only logs that are newer than a relative duration. You can specify the duration in seconds (*Ns*), minutes (*Nm*), or hours (*Nh*).

For example, to view the log files for the scheduler component that are newer than 3 hours, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-scheduler/ {print $1;exit}') -n baas -f --since=3h
```

Troubleshooting issues with snapshot backup jobs

If a snapshot backup operation is unsuccessful, you can take a series of actions to diagnose the problem.

Complete the following steps to troubleshoot snapshot backup problems:

1. Ensure that the Kubernetes Backup Support log files are available. For instructions about viewing the log files, see [“Viewing log files”](#) on page 379.
2. Verify whether the Kubernetes Backup Support scheduler is sending snapshot create requests for the affected volume. Open the scheduler log file and look for the following text:

```
Scheduled a snapshot for volume volumename
```

If the scheduler is not sending snapshot requests, investigate and resolve any scheduler issues.

3. If the scheduler is sending the snapshot request, check the `baas-transaction-manager` container log in the `baas-transaction-manager` pod. In the log file, look for the text `createsnapshot` and `checksnapshot`, and see whether the `createsnapshot` or `checksnapshot` URL contains the volume name as well. Search for text that is similar to the following example:

```
/checksnapshot?requestname=test:test-pvc-415&volumename=test:test-pvc-415&snapshotname=test:test-pvc-415-1569883237
```

4. Look for the outputs from the `/checksnapshot` API call. If you find an exception, review the exception log to help resolve the issue.
5. If no exceptions are found in Step 3, check the `baas-transaction-manager-worker` container logs in the `baas-transaction-manager` pod. Look for the `createSnapshot` job that is processing with the persistent volume claim (PVC) name of the volume that is being backed up. These logs show multiple running processes. Identify the name of the worker in the log files and follow the logs for the worker to determine whether there is an exception.

In the following example of the transaction manager worker `createsnapshot` log, `ForkPoolWorker-28` is the worker:

```
[2019-09-17 02:27:26,362: DEBUG/ForkPoolWorker-28] Received Create: 1568687246  
[2019-09-17 02:27:26,362: INFO/ForkPoolWorker-28] bp.createsnapshot()  
requestname=default:demo-pvc-1
```

You might find the following exceptions in the createsnapshot log:

Table 38. Possible snapshot backup exceptions	
Exception	Action
The snapshot does not exist. The snapshot might not be created properly.	Run the following command to see whether the snapshot was created correctly: <pre>kubectl describe volumesnapshots snapshotname -n namespace</pre>
The deployment does not exist. The data mover might not be created properly.	For more information about the issue, run the following command: <pre>kubectl describe deploy baas-datamover -n baas</pre>
The data mover registration might be failing in IBM Spectrum Protect Plus.	To determine whether the data mover registration failed, look for the text <code>Registering DM with port</code> in the logs. Verify whether an error occurred that is related to the data mover registration. If an error exists with the registration, try deleting the data mover deployment, network policy, and service to help resolve the issue. Look for a job the job with the following naming convention in IBM Spectrum Protect Plus: <pre>kubernetesvol_internalID_namespace:volumename</pre> where <i>internalID</i> is the same as the <i>instanceid</i> value that is generated by Kubernetes Backup Support. If you see this type of job, go to the next step to troubleshoot issues in IBM Spectrum Protect Plus.

6. Troubleshoot IBM Spectrum Protect Plus issues by taking the following actions:

- In the IBM Spectrum Protect Plus user interface, verify whether any inventory jobs are hung that are preventing all other jobs from being recorded in IBM Spectrum Protect Plus.
- Look for the hung job in the list of running jobs or in the job history. Look for job names with the following naming convention:

```
kubernetesvol_internalID_namespace:volumename
```

where the *internalID* value is the same as the *instanceid* value that is generated by Kubernetes Backup Support. The value of the *instanceid* is returned in the output of the **kubectl describe** command that shows the status of a backup or restore job. For more information, see [“Viewing the status of backup and restore jobs”](#) on page 329.

- Check the job logs and resolve any reported issues.

Troubleshooting issues with copy backup jobs

If a copy backup job is unsuccessful, you can take a series of actions to diagnose the problem.

Complete the following steps to troubleshoot copy backup problems:

1. Ensure that the Kubernetes Backup Support log files are available. For instructions about viewing the log files, see [“Viewing log files”](#) on page 379.
2. Verify whether the Kubernetes Backup Support scheduler is sending copy backup requests for the volume. Open the scheduler log file and look for the following text:

```
Scheduled a copy backup for volume volumename
```

If the scheduler is not sending copy backup requests, investigate and resolve the scheduler issues.

3. If the scheduler is sending the snapshot request, check the `baas-transaction-manager` container log in the `baas-transaction-manager` pod. In the log file, look for the text `createcopybackup` and `checkcopybackup`, and see whether the `createcopybackup` or `checkcopybackup` URL contains the volume name as well. Search for text that is similar to the following example:

```
/checkcopybackup?requestname=test:test-pvc-415&volumename=test:test-pvc-415&copybackupname=test:test-pvc-415-1569883237
```

4. Look for the outputs from the `/checkcopybackup` API call. If you find an exception, review the exception log to help resolve the issue.
5. If no exceptions are found in Step 3, check the `baas-transaction-manager-worker` container logs in the `baas-transaction-manager` pod. Look for the `createCopyBackup` job that is processing with the PVC name of the volume that is being backed up. These logs show multiple running processes. Identify the name of the worker in the log files and follow the logs for the worker to determine whether there is an exception.

In the following example of the transaction manager worker `createcopybackup` log, `ForkPoolWorker-28` is the worker:

```
[2019-09-17 02:27:26,362: DEBUG/ForkPoolWorker-28] Received Create: 1568687246
```

You might find the following exceptions in the `createcopybackup` log:

Table 39. Possible copy backup exceptions	
Exception	Action
<p>The snapshot does not exist.</p> <p>The snapshot might not be created properly.</p>	<p>Run the following command to see whether the snapshot was created correctly:</p> <pre>kubect1 describe volumesnapshots <i>snapshotname</i> -n <i>namespace</i></pre>
<p>The deployment does not exist.</p> <p>The data mover might not be created properly.</p>	<p>For more information about the issue, get the data mover name from the error message and run the following command:</p> <pre>kubect1 describe deploy baas-datamover-<i>ipaddress-instanceid</i> -n <i>namespace</i></pre> <p>The <i>instanceid</i> value is returned in the output of the kubect1 describe command that shows the status of a backup or restore job. For more information, see “Viewing the status of backup and restore jobs” on page 329.</p>

Table 39. Possible copy backup exceptions (continued)

Exception	Action
The data mover registration might be failing in IBM Spectrum Protect Plus.	<p>To determine whether the data mover registration failed, look for the text <code>Registering DM with port</code> in the logs. Verify whether an error occurred that is related to the data mover registration.</p> <p>If an error exists with the registration, try deleting the data mover deployment, network policy, and service to help resolve the issue.</p> <p>Look for a job the job with the following naming convention in IBM Spectrum Protect Plus:</p> <pre>kubernetesvol_internalID_namespace:volumename</pre> <p>where <i>internalID</i> is the same as the <i>instanceid</i> value that is generated by Kubernetes Backup Support.</p> <p>If you see this type of job, go to the next step to troubleshoot issues in IBM Spectrum Protect Plus.</p>

6. Troubleshoot IBM Spectrum Protect Plus issues by taking the following actions:

- a. In the IBM Spectrum Protect Plus user interface, verify whether any inventory jobs are hung that are preventing all other jobs from being recorded in IBM Spectrum Protect Plus.
- b. Look for the hung job in the list of running jobs or in the job history. Look for job names with the following naming convention:

```
kubernetesvol_internalID_namespace:volumename
```

where the *internalID* value is the same as the *instanceid* value that is generated by Kubernetes Backup Support. The value of the *instanceid* is returned in the output of the **kubect1 describe** command that shows the status of a backup or restore job. For more information, see [“Viewing the status of backup and restore jobs” on page 329](#).

- c. Check the job logs and resolve any reported issues.

Troubleshooting restore jobs

If a restore job is unsuccessful, you can take the following actions to diagnose the problem.

Complete the following steps to troubleshoot restore job problems:

1. Ensure that the Kubernetes Backup Support log files are available. For instructions about viewing the log files, see [“Viewing log files” on page 379](#).
2. Check the controller log file to see whether the restore request was invalidated due to a data issue.

Look for the restore request name in the controller log. Any errors are listed by the request name.

If there are no errors, the `/checkrestore` calls are shown in the transaction manager logs. If errors exist, create the restore request correctly.

If the restore request is created correctly, the scheduler receives the request.
3. Verify whether the volume to be restored has an ongoing snapshot or copy backup in process. You can accomplish this task by looking at the `/checkcopybackup` or `/checksnapshot` calls that are being made in the transaction manager log for this volume.

If there are no exceptions, the `/restorebackup` and `/checkrestorebackup` calls are shown in the transaction manager logs for this volume. Any exceptions in the `/restorebackup` calls are displayed with a return code in the log.

Restriction: If you are troubleshooting a fast restore job, skip the following steps in this procedure.

4. Resolve any exceptions by following the instructions in the following table.

Table 40. Possible restore exceptions (skip for fast restores)	
Symptom	Action
<p>The deployment does not exist.</p> <p>The data mover might not be created properly.</p>	<p>For more information about the issue, get the data mover name from the error message and run the following command:</p> <pre>kubect1 describe deploy baas-datamover- ipaddress-instanceid -n namespace</pre> <p>The <i>instanceid</i> value is returned in the output of the kubect1 describe command that shows the status of a backup or restore job. For more information, see “Viewing the status of backup and restore jobs” on page 329.</p>
<p>The data mover registration might be failing in IBM Spectrum Protect Plus.</p>	<p>To determine whether the data mover registration failed, look for the text <code>Registering DM with port</code> in the logs. Verify whether an error occurred that is related to the data mover registration.</p> <p>If an error exists with the registration, try deleting the data mover deployment, network policy, and service to help resolve the issue.</p> <p>Look for a job the job with the following naming convention in IBM Spectrum Protect Plus:</p> <pre>kubernetesvol_internalID_namespace:volumena me</pre> <p>where <i>internalID</i> is the same as the <i>instanceid</i> value that is generated by Kubernetes Backup Support.</p> <p>If you see this type of job, go to the next step to troubleshoot issues in IBM Spectrum Protect Plus.</p>

5. Troubleshoot IBM Spectrum Protect Plus issues by taking the following actions. Skip this step for fast restores.

- a. In the IBM Spectrum Protect Plus user interface, verify whether any inventory jobs that are hung are preventing all other jobs from being recorded in IBM Spectrum Protect Plus.
- b. Look for the IBM Spectrum Protect Plus job in the list of running jobs or in the job history. Look for job names that have the following naming convention:

```
onDemandRestore_timestamp
```

To verify that the job applies to this specific volume, view the IBM Spectrum Protect Plus job logs and verify the *internalID* and *dbname* values that are associated with the volume's data.

- c. Resolve any issues that are reported in the relevant job logs.

Related tasks

[“Setting the trace level of log files” on page 386](#)

You can set the trace level of local log files to help troubleshoot issues that you might encounter in Kubernetes Backup Support.

Related reference

[“Troubleshooting quick reference” on page 377](#)

Solutions to basic Kubernetes Backup Support problems are provided.

Collecting Kubernetes Backup Support log files for troubleshooting

You can generate debugging log files in the Kubernetes environment to troubleshoot the deployment of Kubernetes Backup Support and Kubernetes Backup Support operations on the IBM Spectrum Protect Plus server.

About this task

All logs are collected in the `/tmp` directory on the local system and packaged into a `tar.gz` archive file. The archive file is typically named `baas_debug_logs_timestamp.tar.gz`.

Procedure

Use one of the following methods to collect logs for troubleshooting:

- To collect only Kubernetes logs for debugging purposes, issue the following command:

```
./baas_install.sh -l
```

This command collects debugging logs for the Kubernetes Backup Support deployment that is specified by the parameters in the `baas_config.cfg`. The current state information and logs of the Kubernetes Backup Support components in the Kubernetes cluster are collected. The logs are structured based on the Kubernetes basic logging architecture. For more information, see [Basic logging in Kubernetes](#).

- To collect the log package that includes the debugging logs for the Kubernetes Backup Support deployment and IBM Spectrum Protect Plus server, issue the following command:

```
./baas_install.sh -l -x
```

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.

Related tasks

[“Setting the trace level of log files” on page 386](#)

You can set the trace level of local log files to help troubleshoot issues that you might encounter in Kubernetes Backup Support.

Related reference

[“Troubleshooting quick reference” on page 377](#)

Solutions to basic Kubernetes Backup Support problems are provided.

[“Troubleshooting Kubernetes Backup Support operations” on page 379](#)

Troubleshooting procedures are available to help you diagnose and resolve Kubernetes Backup Support issues.

Setting the trace level of log files

You can set the trace level of local log files to help troubleshoot issues that you might encounter in Kubernetes Backup Support.

About this task

You can set the trace levels to troubleshoot issues with the Kubernetes Backup Support transaction manager, controller, and scheduler components. To set the trace level, you must update the `baas_config.cfg` configuration file and then update the Kubernetes Backup Support deployment.

The data mover component is not affected by this setting.

Procedure

To set the trace level for the Kubernetes Backup Support transaction manager, controller, and scheduler log files, complete the following steps in the Kubernetes environment:

1. Log in to the operating system on the master node of the Kubernetes cluster that is used as the installation node.
2. Go to the directory where the `installer-10.1.5.tar.gz` installation package was unpacked.
3. Go to the `installer` directory by issuing the following command:

```
cd installer
```

4. Edit the `baas_config.cfg` file with a text editor and modify the value for the **PRODUCT_LOGLEVEL** parameter.

The following trace options are available:

INFO

Display all user messages in the transaction manager, controller, and scheduler log files, including information, warning, and error messages. This value is the default.

WARNING

Display warning and error messages in the transaction manager, controller, and scheduler log files.

ERROR

Display only error messages in the transaction manager, controller, and scheduler log files.

DEBUG

Display debugging-level messages in the transaction manager, controller, and scheduler log files.

For example, to set the trace level to debugging mode, set the **PRODUCT_LOGLEVEL** parameter as follows:

```
PRODUCT_LOGLEVEL="DEBUG"
```

5. Update the Kubernetes Backup Support deployment by issuing the following command:

```
./baas_install.sh -u
```

When prompted, enter yes to continue.

6. Optional: To verify the status of the update, issue the following command:

```
./baas_install.sh -s
```

Tip: Alternatively, verify the status of the update by using the `./helm status baas` command.

What to do next

You can collect Kubernetes Backup Support log files for troubleshooting or use a visualization tool such as Kibana to view and query data in the transaction manager, controller, and scheduler log files. For instructions, see:

- [“Collecting Kubernetes Backup Support log files for troubleshooting” on page 385](#)
- [“Viewing trace logs for Kubernetes Backup Support” on page 387](#)

Viewing trace logs for Kubernetes Backup Support

You can optionally use the Elasticsearch, Fluentd, and Kibana (EFK) stack to view and analyze trace logs that are produced by Kubernetes Backup Support.

Elasticsearch is a distributed full-text search engine. Fluentd is a tool that collects logs from cluster nodes and sends the logs to the Elasticsearch engine. Kibana is a visualization tool for Elasticsearch with a web user interface and development tool that is used for querying data.

Before you begin

Complete the following steps:

1. Deploy the EFK stack to your Kubernetes cluster:
 - a. Deploy the Elasticsearch search engine. For instructions, see [Installing Elasticsearch](#).
 - b. Deploy the Fluentd log collector on each cluster node. For instructions, see the [Fluentd documentation](#).
 - c. Deploy the Kibana visualization tool. For instructions, see the [Kibana Guide](#).
2. Complete the EFK stack deployment by adding a logstash index in Kibana:
 - a. Access the Kibana user interface by opening a web browser and entering the URL of the computer where Kibana is running and specify the port number. For example, specify one of the following URLs in your web browser:

```
https://localhost:5601
```

or

```
http://your_domain.com:5601
```

where *your_domain* specifies the domain name for the computer.
 - b. If you are prompted with options to explore data, select **Explore on my own**.
 - c. Click the **Discover > Create Index Pattern** and create the logstash-* index pattern.

About this task

When you use the EFK stack, the logs from all container components are merged and shown in the same view. Any logs for stopped pods are preserved in Elasticsearch persistent data storage. You can apply filters to display specific errors or messages. You can also apply a time filter to show events that occurred in a specific time period.

In addition to error and debugging messages, you can view trace logs for the following Kubernetes Backup Support components:

- Transaction manager
- Controller
- Scheduler

Procedure

To view transaction logs for Kubernetes Backup Support, complete the following steps:

1. Open the Kibana user interface and click the **Discover** icon.
2. Click the `logstash-*` index.
3. To view logs for Kubernetes Backup Support, add a filter by taking the following actions:
 - a) Click **Add filter** and specify the following filter values:
 - Field: `kubernetes.container_image`
 - Operator: `is`
 - Value: `baas-`
 - b) Enter a name for the search and click **Save**.
 The trace logs for the `baas-transaction-manager`, `baas-controller`, and `baas-scheduler` containers are displayed.
4. You can create additional filters to show more granular views of Kubernetes Backup Support trace logs.

Table 41. Filters for viewing Kubernetes Backup Support trace logs		
Type of data to show	Filter 1	Filter 2
Transaction manager logs	<code>kubernetes.container_image is baas-transaction-manager</code>	None
Controller logs	<code>kubernetes.container_image is baas-controller</code>	None
Scheduler logs	<code>kubernetes.container_image is baas-scheduler</code>	None
Error messages	<code>kubernetes.container_image is baas-</code>	<code>log is ERROR</code>
Debugging messages	<code>kubernetes.container_image is baas-</code>	<code>log is DEBUG</code>

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.

Table 42. Messages prefixes by 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
CTGGR	IBM Spectrum Protect Plus for Microsoft Office 365

For a list of all messages, see IBM Knowledge Center [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](http://www.w3.org/TR/wai-aria/) (www.w3.org/TR/wai-aria/), to ensure compliance with [US Section 508](http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and [Web Content Accessibility Guidelines \(WCAG\) 2.0](http://www.w3.org/TR/WCAG20/) (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the [Accessibility](http://www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility) section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

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

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

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

Vendor software

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

Related accessibility information

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

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

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

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

Trademarks

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

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

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

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

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

Rights

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

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

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

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

Privacy policy considerations

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

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

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

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

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.
See the [IBM Spectrum Protect glossary](#).

Index

A

- Access control
 - MongoDB [253](#)
- accessibility features [393](#)
- ad hoc jobs
 - creating [349](#)
- Add Db2 partitions [192](#)
- adding
 - Hyper-V servers [174](#)
 - identities [373](#)
 - LDAP server [127](#)
 - Oracle application servers [282](#)
 - sites [124](#)
 - SMTP server [128](#)
 - SQL Server application servers [294](#)
 - vCenter Server instances [151](#)
 - virtual disks to a vCenter virtual machine [141](#)
 - vSnap servers [73](#)
- Adding Db2 [192](#)
- Adding MongoDB [254](#)
- Administrative Console, logging on to [135](#)
- Advanced backup options [78](#)
- application server
 - Db2 [189](#)

B

- backing up
 - container data [322](#)
- Backing up
 - Db2 [196](#)
- backing up container data
 - on demand [324](#)
 - scheduling [322](#)
- backup encryption
 - Kubernetes Backup Support [322](#)
- backup jobs
 - creating
 - Hyper-V [176](#)
 - IBM Spectrum Protect Plus [339](#)
 - Oracle [284](#)
 - SQL Server [296](#)
 - VMware [155](#)
 - excluding VMDKs from [159](#)
 - rerunning
 - on demand [349](#)
 - starting
 - on demand [345](#)
 - on schedule [145](#)
- backup policies, *See* SLA policies
- backup storage
 - advanced options, managing [78](#)
 - storage options, managing disks [75](#)
 - storage options, managing partners [77](#)
- backup storage server
 - storage options, managing [77](#)

- backup types
 - Kubernetes Backup Support [308](#)
- Beta program
 - advantages [xiii](#)
 - overview [xiii](#)

C

- certificate
 - adding [122](#)
 - deleting [122](#)
- cloud provider
 - deleting [116](#)
 - editing [116](#)
- cloud server
 - adding a Microsoft azure cloud resource [114](#)
 - adding an Amazon S3 [111](#)
 - adding an IBM Cloud Object Storage resource [112](#)
 - adding an s3 compatible cloud resource [115](#)
- collecting debugging log filesKubernetes Backup Support [385](#)
- completely uninstalling
 - Kubernetes Backup Support [321](#)
- configuring
 - Kubernetes Backup Support [315](#)
- Configuring backup storage
 - storage options, adding disks [76](#)
- copy backup
 - Kubernetes Backup Support [322](#)
- copy restore
 - container data [326](#)
- creating
 - reports [360](#)
 - resource groups [364](#)
 - roles [368](#)
 - SLA policies [145](#)
 - users
 - individual [371](#)
 - LDAP group [371](#)
 - VADP proxies [161](#)
- customizing
 - global
 - preferences [133](#)

D

- Db2
 - system requirements [36](#)
- Db2 log backup [201](#)
- deleting
 - identities [374](#)
 - jobs [348](#)
 - LDAP server [129](#)
 - resource groups [367](#)
 - roles [371](#)
 - sites [126](#)
 - SLA demo [149](#)

- deleting (*continued*)
 - SLA policies [149](#)
 - SMTP server [129](#)
 - users [373](#)
- deleting backups
 - Kubernetes Backup Support [334](#)
- deploying
 - Kubernetes Backup Support [315](#)
- deployment log files [Kubernetes Backup Support 385](#)
- destroying backups [Kubernetes Backup Support 334](#)
- Detecting
 - Db2 [194](#)
- disability [393](#)

E

- early availability updates, obtaining and applying [109](#)
- editing
 - identities [373](#)
 - jobs and job schedules [347](#)
 - LDAP server [129](#)
 - resource groups [366](#)
 - roles [370](#)
 - settings [129](#)
 - sites [125](#)
 - SLA policies [149](#)
 - SMTP server [129](#)
 - users [372](#)
- efix [109](#)
- enable tracing
 - Kubernetes Backup Support [386](#)
- Exchange Server
 - system requirements [31](#)
- expire job session [340](#)

F

- fast restore
 - container data [326](#)
- fenced network, creating [170](#)
- files
 - restoring [185](#)
 - searching for [391](#)
- Finding Db2 [194](#)
- firewalls [63](#)

G

- global preferences
 - configuring [130](#)
 - schema driven [133](#)

H

- Hyper-V
 - adding [174](#)
 - backup job, creating [176](#)
 - installing on virtual appliance [60](#)
 - restore job, creating [179](#)
 - servers
 - detecting resources for [175](#)
 - enabling WinRM [175](#)
 - testing connection to [176](#)

- Hyper-V (*continued*)
 - virtual appliance
 - accessing [139](#)

I

- IBM Knowledge Center [ix](#)
- IBM spectrum protect server
 - adding a repository server [119](#)
- identities
 - adding [373](#)
 - deleting [374](#)
 - editing [373](#)
- installing
 - download packages, obtaining [58](#)
 - Kubernetes Backup Support [313](#), [315](#)
 - virtual appliance
 - on Hyper-V [60](#)
 - on VMware [58](#)
 - vSnap servers
 - Hyper-V environment [69](#)
 - physical environment [67](#)
 - VMware environment [68](#)
- iSCSI utilities
 - installing [65](#)

J

- jobs
 - canceling [348](#)
 - concurrent [344](#)
 - creating [344](#)
 - deleting [348](#)
 - editing [347](#)
 - logs, downloading [346](#)
 - names of [343](#)
 - pausing [347](#)
 - releasing [347](#)
 - rerunning [349](#)
 - schedules, editing [347](#)
 - starting
 - on demand [345](#)
 - on schedule [145](#)
 - types of [343](#)
 - viewing [346](#)
- Jobs and Operations [343](#)

K

- key
 - adding [121](#), [123](#)
 - deleting [122](#), [123](#)
- keyboard [393](#)
- keys [121](#)
- Knowledge Center [ix](#)
- Kubernetes Backup Support
 - baas requests [310](#)
 - backing up container data [322](#)
 - backup and restore types [308](#)
 - backup status [332](#)
 - cascading actions [313](#)
 - collecting debugging log files [385](#)
 - complete uninstall [321](#)

Kubernetes Backup Support *(continued)*

- configuration file [315](#)
- copy backup [322](#)
- copy restore [326](#)
- deleting backups [334](#)
- deployment logs [385](#)
- destroy request [334](#)
- displaying log files [379](#)
- enable tracing [386](#)
- enable VolumeSnapshotDataSource feature [313](#)
- encrypting backups [322](#)
- encryption [311](#)
- fast restore [326](#)
- installing [313](#), [315](#)
- managing jobs [329](#)
- monitoring jobs [336](#)
- multitenancy [311](#)
- overview [307](#)
- pausing scheduled backups [333](#)
- prerequisites [313](#)
- request types [310](#)
- restore status [332](#)
- restoring data [326](#)
- resuming scheduled backups [334](#)
- running reports [336](#)
- scheduling backups [322](#)
- security [311](#)
- setting trace levels [386](#)
- SLA policies [309](#), [322](#)
- snapshot backup [322](#), [324](#)
- system requirements [54](#)
- troubleshooting [377](#)
- troubleshooting backup jobs [379](#)
- troubleshooting restore jobs [379](#)
- uninstalling [320](#)
- user roles [309](#)
- verifying metrics server [313](#)
- viewing backup history [337](#)
- viewing backup status [329](#)
- viewing job logs [336](#)
- viewing restore status [329](#)
- viewing trace logs [387](#)

L

LDAP

- group, creating a user account for [371](#)
- server
 - adding [127](#)
 - deleting [129](#)
 - settings, editing [129](#)

Linux-based vCenter virtual appliance, backing up [160](#)

Log archiving

- Db2 [201](#)

logs

- audit
 - downloading [361](#)
 - viewing [361](#)
- system
 - downloading [361](#), [377](#)
 - viewing [361](#), [377](#)

M

- managing jobs
 - container backups and restores [329](#)
- message
 - prefixes [389](#)
- messages [389](#)
- MongoDB
 - system requirements [39](#)
- MongoDB application server [252](#)
- monitoring
 - container backup jobs [336](#)
- multitenancy
 - Kubernetes Backup Support [307](#), [311](#)

N

- network
 - testing [139](#), [140](#)
- New in IBM Spectrum Protect Plus Version 10.1.5 [xi](#)

O

- Object Storage
 - Amazon S3 [111](#)
- offline updates [103](#)
- on-demand backup
 - containers [324](#)
- online updates [103](#)
- Ops Manager
 - MongoDB [256](#)
- Oracle
 - application servers
 - adding [282](#)
 - detecting resources for [284](#)
 - testing connection to [284](#)
 - backup job, creating [284](#)
 - multithreaded databases [282](#)
 - restore job, creating [287](#)
 - system requirements [44](#)
- overview
 - Kubernetes Backup Support [307](#)

P

- pausing scheduled backups
 - Kubernetes Backup Support [333](#)
- preferences
 - global
 - configuring [130](#)
 - schema driven [133](#)
- prerequisites
 - Db2 [189](#)
 - Kubernetes Backup Support [313](#)
 - MongoDB [252](#)
- Prerequisites
 - MongoDB [253](#)
- publications [ix](#)

Q

- quick start [89](#)

R

- RBAC
 - MongoDB [253](#)
- repair vSnap [83](#)
- Replication partners [77](#)
- reports
 - custom, creating [360](#)
 - running
 - on demand [359](#)
 - on schedule [360](#)
 - running VM [357](#)
 - types of
 - backup storage utilization [353](#)
 - protection [354](#)
 - system [356](#)
- repository server provider
 - deleting [121](#)
 - editing [121](#)
- request types
 - Kubernetes Backup Support [310](#)
- rerunning
 - jobs
 - on demand [349](#)
- resource groups
 - creating [364](#)
 - deleting [367](#)
 - editing [366](#)
 - types of [365](#)
- restore jobs
 - creating
 - Hyper-V [179](#)
 - IBM Spectrum Protect Plus [339](#)
 - Oracle [287](#)
 - SQL Server [300](#)
 - VMware [164](#)
 - running
 - Hyper-V [179](#)
 - Oracle [287](#)
 - SQL Server [300](#)
 - VMware [164](#)
- restore points, deleting [341](#)
- restore points, managing [340](#)
- restore types
 - Kubernetes Backup Support [308](#)
- Restoring
 - Db2 [202](#), [207](#), [210](#)
- restoring container data
 - Kubernetes Backup Support [326](#)
- Restoring Db2
 - Alternate instance [210](#)
 - Original instance [207](#)
- resuming scheduled backups
 - Kubernetes Backup Support [334](#)
- roles
 - creating [368](#)
 - deleting [371](#)
 - editing [370](#)
 - permission types [369](#)
- running reports
 - container backup jobs [336](#)

S

- Schedule jobs
 - Backup [198](#), [218](#), [260](#)
- scheduling backups
 - Kubernetes Backup Support [322](#)
- scripts for backup and restore operations
 - uploading [350](#), [351](#)
- security features
 - Kubernetes Backup Support [311](#)
- service level agreement, *See* SLA policies
- service level agreements
 - Kubernetes Backup Support [309](#)
- Setting Db2
 - SLA options [200](#)
- setting trace levels
 - Kubernetes Backup Support [386](#)
- sites
 - adding [124](#)
 - deleting [126](#)
 - editing [125](#)
 - throttling [124](#), [125](#)
- SLA [198](#), [218](#), [260](#)
- SLA options
 - Db2 [200](#)
- SLA policies
 - adding [145](#)
 - deleting [149](#)
 - editing [149](#)
 - Kubernetes Backup Support [309](#), [322](#)
- SMTP
 - server
 - adding [128](#)
 - deleting [129](#)
 - settings, editing [129](#)
- snapshot backup
 - containers [324](#)
 - Kubernetes Backup Support [322](#)
- snapshot retention [340](#)
- sponsor user program
 - advantages [xiii](#)
 - overview [xiii](#)
- SQL Server
 - application servers
 - adding [294](#)
 - detecting resources for [296](#)
 - testing connection to [296](#)
 - backup job, creating [296](#)
 - requirements for data protection [293](#)
 - restore job, creating [300](#)
 - system requirements [49](#)
- SSL certificate, uploading
 - from administrative console [137](#)
 - from command line [138](#)
- starting
 - IBM Spectrum Protect Plus [91](#)
 - jobs
 - on demand [345](#)
 - on schedule [145](#)
- system requirements
 - components [11](#)
 - Db2 [36](#)
 - Exchange Server [31](#)
 - file index and restore [27](#)

- system requirements (*continued*)
 - hypervisors [26](#)
 - Kubernetes Backup Support [54](#)
 - MongoDB [39](#)
 - Oracle [44](#)
 - SQL Server [49](#)
- system(
 - logs [361](#)

T

- Testing connection
 - Db2 [195](#)
- time zone, setting [136](#)
- troubleshooting
 - displaying Kubernetes Backup Support logs [379](#)
 - Kubernetes Backup Support [377](#)
 - Kubernetes Backup Support operations [379](#)

U

- uninstalling
 - Kubernetes Backup Support [320](#)
- user access [5](#), [363](#)
- user roles
 - Kubernetes Backup Support [309](#)
- users
 - deleting [373](#)
 - editing [372](#)
 - individual, creating [371](#)
 - LDAP group, creating [371](#)
 - resource groups
 - creating [364](#)
 - deleting [367](#)
 - editing [366](#)
 - types of [365](#)
 - roles
 - creating [368](#)
 - deleting [371](#)
 - editing [370](#)
 - permission types [369](#)

V

- VADP proxies
 - creating [161](#)
 - options, setting [163](#)
 - uninstalling [164](#)
 - updating [108](#)
- verifying metrics server
 - Kubernetes Backup Support [313](#)
- viewing backup history
 - container backups [337](#)
- viewing backup status
 - Kubernetes Backup Support [329](#), [332](#)
- viewing job logs
 - container backups [336](#)
- viewing restore status
 - Kubernetes Backup Support [329](#), [332](#)
- viewing trace logs
 - Kubernetes Backup Support [387](#)
- virtual appliance
 - accessing

- virtual appliance (*continued*)
 - accessing (*continued*)
 - in Hyper-V [139](#)
 - in VMware [138](#)
 - adding a disk to [141](#)
 - adding storage capacity [142](#)
 - installing
 - on Hyper-V [60](#)
 - on VMware [58](#)
- Virtual appliance
 - updating [103](#)
- VMware
 - backup job, creating [155](#)
 - backup job, excluding VMDKs from SLA policy [159](#)
 - installing on virtual appliance [58](#)
 - restore job
 - creating a fenced network [170](#)
 - restore job, creating [164](#)
 - vCenter Server instances
 - adding [151](#)
 - vCenter Server, detecting resources [155](#)
 - vCenter Server, testing connection to [155](#)
 - virtual appliance
 - accessing [138](#)
 - virtual machine privileges, required [152](#)
- VolumeSnapshotDataSource feature
 - Kubernetes Backup Support [313](#)
- vSnap
 - updating [107](#)
- vSnap recovery [83](#)
- vSnap server
 - administering
 - kernel headers
 - kernel tools [87](#)
 - network administration [87](#)
 - storage administration [84](#)
 - change throughput [82](#)
 - deleting [74](#)
 - editing [74](#)
 - initializing
 - advanced [81](#)
 - simple [80](#)
 - replication partnership, establishing [82](#)
 - storage pools, expanding [81](#)
- vSnap servers
 - adding [73](#)
 - installing
 - Hyper-V environment [69](#)
 - physical environment [67](#)
 - VMware environment [68](#)
 - uninstalling [70](#)

W

- WinRM, enabling for connection to Hyper-V servers [175](#)

Y

- YAML files
 - Kubernetes Backup Support [310](#)



Product Number: 5737-F11

Printed in USA