

IBM Spectrum Protect Plus
Version 10.1.2

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 187.](#)

This edition applies to version 10, release 1, modification 2 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, 2018.**

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

Contents

What's new in V10.1.2.....	vii
Chapter 1. Product overview.....	1
Product components.....	1
Product dashboard.....	3
Product deployment roadmap.....	4
Alerts.....	5
Role-based access control.....	5
Replicate backup-storage data.....	6
Offload to IBM Spectrum Protect.....	6
Data Protection for VMware deployment checklist.....	8
IBM Spectrum Protect Plus on IBM® Cloud.....	11
Chapter 2. Installing IBM Spectrum Protect Plus.....	13
System requirements	13
Component requirements	13
Hypervisor requirements	21
File indexing and restore requirements.....	21
Db2 requirements.....	24
Oracle requirements.....	27
Microsoft SQL Server requirements.....	30
Obtaining the IBM Spectrum Protect Plus installation package.....	33
Installing IBM Spectrum Protect Plus as a VMware virtual appliance.....	34
Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance.....	36
Assigning a static IP address.....	37
Uploading the product key.....	38
Chapter 3. Installing and configuring vSnap servers.....	39
Installing vSnap servers.....	39
Installing a physical vSnap server.....	39
Installing a virtual vSnap server in a VMware environment.....	40
Installing a virtual vSnap server in a Hyper-V environment.....	41
Managing vSnap servers.....	42
Adding a vSnap server as a backup storage provider.....	42
Initializing the vSnap server.....	44
Setting vSnap storage options.....	45
Expanding a vSnap storage pool.....	45
Establishing a replication partnership for vSnap servers.....	46
vSnap server administration reference	46
Storage management.....	47
Network management.....	49
Uninstalling a vSnap server.....	50
Chapter 4. Getting off to a quick start.....	51
Start IBM Spectrum Protect Plus.....	52
Create backup policies.....	52
Create a user account for the application administrator.....	54
Add resources to protect.....	56
Add resources to a job definition.....	58
Start a job.....	59
Run a report.....	60

Chapter 5. Updating IBM Spectrum Protect Plus components.....	61
Updating the IBM Spectrum Protect Plus virtual appliance.....	61
Updating vSnap servers.....	63
Updating vSnap servers in a Red Hat Enterprise Linux environment.....	64
Updating VADP proxies.....	64
Chapter 6. Managing SLA policies for backup operations.....	67
Creating an SLA policy.....	67
Editing an SLA policy.....	69
Deleting an SLA policy.....	69
Chapter 7. Protecting hypervisors.....	71
VMware.....	71
Adding a vCenter Server instance.....	71
Backing up VMware data.....	75
Managing VADP backup proxies.....	79
Restoring VMware data.....	82
Hyper-V.....	90
Adding a Hyper-V server.....	90
Backing up Hyper-V data.....	92
Restoring Hyper-V data.....	95
Restoring files.....	98
Chapter 8. Protecting applications.....	101
Db2.....	101
Prerequisites for Db2.....	101
Adding a Db2 application server.....	103
Backing up Db2 data.....	106
Restoring Db2 data	110
SQL Server.....	122
Adding a SQL Server application server.....	123
Backing up SQL Server data.....	125
Restoring SQL Server data.....	127
Oracle.....	131
Adding an Oracle application server.....	131
Backing up Oracle data.....	133
Restoring Oracle data.....	135
Chapter 9. Protecting IBM Spectrum Protect Plus.....	141
Backing up the application.....	141
Restoring the application.....	141
Managing restore points.....	142
Chapter 10. Managing jobs.....	143
Job types.....	143
Starting jobs.....	144
Pausing and resuming jobs.....	144
Canceling jobs.....	144
Configuring scripts for backup and restore operations.....	144
Uploading a script.....	145
Adding a script to a server.....	145
Chapter 11. Configuring and maintaining the IBM Spectrum Protect Plus system environment.....	147
Managing sites.....	147

Adding a site.....	147
Editing a site.....	147
Deleting a site.....	148
Managing LDAP and SMTP servers.....	148
Adding an SMTP server.....	148
Adding an LDAP server.....	149
Logging on to the administrative console.....	150
Setting the time zone.....	151
Uploading an SSL certificate from the administrative console.....	152
Uploading an SSL certificate from the command line.....	152
Logging on to the virtual appliance.....	153
Accessing the virtual appliance in VMware.....	153
Accessing the virtual appliance in Hyper-V.....	153
Testing network connectivity.....	154
Running the Service Tool from a command-line interface.....	154
Running the Service Tool remotely.....	155
Adding virtual disks.....	155
Adding a disk to the virtual appliance.....	156
Adding storage capacity from a new disk to the appliance volume.....	156
Chapter 12. Managing reports and logs.....	161
Types of reports.....	161
Backup storage utilization reports.....	161
Protection reports.....	162
System reports.....	163
VM Environment reports.....	164
Report actions.....	165
Running a report.....	165
Creating a custom report.....	165
Scheduling a report.....	166
Collecting and reviewing audit logs for actions.....	166
Chapter 13. Managing user access.....	169
Managing user resource groups.....	170
Creating a resource group.....	170
Editing a resource group.....	172
Deleting a resource group.....	172
Managing roles.....	173
Creating a role.....	174
Editing a role.....	176
Deleting a role.....	176
Managing user accounts.....	177
Creating a user account for an individual user.....	177
Creating a user account for an LDAP group.....	177
Editing a user account credentials.....	178
Deleting a user account.....	178
Managing identities.....	179
Adding an identity.....	179
Editing an identity.....	179
Deleting an identity.....	179
Chapter 14. Troubleshooting.....	181
Collecting log files for troubleshooting.....	181
Appendix A. Search guidelines.....	183
Appendix B. Accessibility.....	185

Notices.....	187
Glossary.....	191
Index.....	193

What's new in Version 10.1.2

IBM Spectrum Protect Plus Version 10.1.2 introduces new features and updates.

For a list of new features and updates in this release and previous Version 10 releases, see [What's new](#).

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

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 your IBM Spectrum Protect environment to offload copies for long-term 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 Blueprint](#).

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 appliance 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.

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 (a data center location) or logical (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:

Alerts icon 	This icon opens the Alerts window. For more information about alerts, see “Alerts” on page 5 .
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.

Sample deployment

[Figure 1](#) 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.

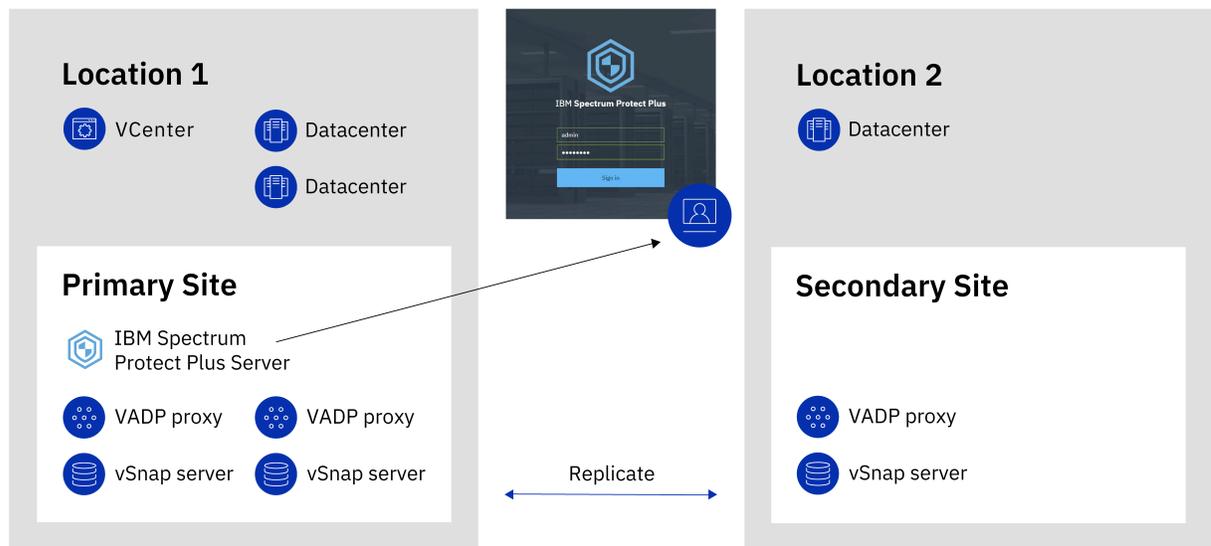


Figure 1: IBM Spectrum Protect Plus deployment across two geographical locations

Product dashboard

The IBM Spectrum Protect Plus dashboard provides the following summary information so that you can monitor the health of your virtual environment.

Job Status

The **Job Status** section displays a summary of job activities over a selected time period. You can select a time period from the **Job status period** field to display job status information from the selected period in the overview and scatter plot. The following information displays in the overview of the **Job Status** section:

Successful

The total number of jobs within the selected period that completed with no warnings or critical errors.

Failed

The total number of jobs within the selected period that failed with critical errors or that wholly failed to complete.

Warning

The total number of jobs within the selected period that partially completed, were skipped, or otherwise resulted in warning errors.

Running

The total number of jobs that are currently running.

Total Jobs

The total number of jobs set to complete processing within the selected time period, including running jobs.

Success Rate

The success rate for jobs over the selected period. The success rate is calculated by using the following formula:

$$100 \times \text{Successful Jobs} / (\text{Total Jobs} - \text{Running Jobs}) = \text{Success Rate}$$

The **Job Status Over Time** section displays the status of job sessions on a scatter plot. You can add or remove the **Failed**, **Warning**, and **Successful** lanes by selecting or clearing the check boxes above the graph. Each circle on the graph represents a collection of jobs that completed during the selected period. Hover your cursor over the circles to view the total number of jobs represented by the circle. Click a circle to display the details of the jobs, including the duration, status, and associated job logs.

Capacity Summary

The **Capacity Summary** section displays the usage and availability of the vSnap server available to IBM Spectrum Protect Plus. Select a time period from the **Storage error period** field to display storage capacity information from the selected period. The following information displays in the overview of the **Capacity Summary** section:

Unavailable Devices

The total number of storage resources that are offline or otherwise unavailable.

Full Devices

The total number of storage resources that have a free capacity of 0%.

Storage resources are listed in the table beneath the overview section. You can expand **Unavailable Devices** and **Full Devices**, and then click a storage resource to view the associated capacity errors. The following fields are displayed:

Time

The time that the error was logged.

Site

The site where the storage resource is located.

Type

The storage resource type

Error

The error message logged in the system at the shown time. You can click **View Log** to review the log at the time the error occurred.

The **System Capacity Breakdown** section displays the usage of your storage resources as well as their deduplication and compressions ratios. The following information is displayed in the **System Capacity Breakdown** section:

Space Available

The total amount of space remaining across all storage resources.

Average Deduplication Ratio

The ratio of the amount of data that is protected compared with the physical space required to store the data after duplicates are removed.

Average Compression Ratio

The ratio of the amount of data that is protected compared with the physical space required to store the data following data compression.

You can click **View backup storage** to view your current storage configuration on the **Backup Storage** pane and resolve capacity issues. From the **Backup Storage** pane, you can add additional vSnap servers or expand the storage pools of your existing vSnap servers. For information about how to install vSnap servers and expand storage pools, see [“Installing vSnap servers” on page 39](#) and [“Expanding a vSnap storage pool” on page 45](#).

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 13 .
Determine how to size, build, and place the components in your IBM Spectrum Protect Plus environment.	See the IBM Spectrum Protect Plus Blueprint .
Install IBM Spectrum Protect Plus.	See Chapter 2, “Installing IBM Spectrum Protect Plus,” on page 13 .

Action	How to
If additional vSnap servers are required to support your environment, install and configure the servers.	See Chapter 3, “Installing and configuring vSnap servers,” on page 39.
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 79.
Complete the basic steps to set up and start using IBM Spectrum Protect Plus.	See Chapter 4, “Getting off to a quick start,” on page 51.

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.

System disk space low

Is displayed when the amount of data consumed by the catalog on an IBM Spectrum Protect Plus data disk exceeds the assigned 5% threshold.

vSnap storage space low

Is displayed when the amount of free disk space on a vSnap server is less than 5%.

System memory low

Is displayed when the amount of memory available to run IBM Spectrum Protect Plus is less than 5%.

System CPU usage high

Is displayed when IBM Spectrum Protect Plus processor usage exceeds the assigned 5% threshold.

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 13, “Managing user access,”](#) on page 169.

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 6, “Managing SLA policies for backup operations,”](#) on page 67.

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. In each SLA policy, in the **Backup Storage Target Site** section, select the primary site. In the **Backup Storage Replication Target Site** section, select the target replication sites.

Offload to IBM Spectrum Protect

Although the vSnap server is the primary backup location for snapshots, you can also offload snapshots to IBM Spectrum Protect.

Prerequisites for offloading snapshots: To use the offload feature, you must have IBM Spectrum Protect server and IBM Spectrum Protect for Virtual Environments: Data Protection for VMware installed and configured. It is assumed that you have an understanding of these applications.

When you use the offload feature in IBM Spectrum Protect Plus, two backups of your data are created. One backup is on the vSnap server for short-term protection, and the other backup is on a IBM Spectrum Protect server for long-term protection.

Enabling the offload feature

The offload 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 6, “Managing SLA policies for backup operations,”](#) on page 67.

The SLA policy determines the method that is used to offload snapshots to IBM Spectrum Protect server. These methods are:

Offload from the hypervisor

The offload operation happens from the hypervisor directly. Incremental backups are supported. This is the default method.

Offload from the vSnap server

The offload operation happens from the vSnap server. Incremental backups are not supported. All backups are full backups. To enable this method, select the **Leverage most recent backup** option for the SLA policy.

Important: Microsoft Hyper-V is not currently supported for offload for either method.

The decision regarding which offload method to choose depends on the use case and environment. Factors to consider when you are deciding which offload method to use include speed, impact on production hypervisor servers, and storage needs.

Consider offloading from the hypervisor if you frequently create long-term recovery points (for example, daily).

Consider offloading from the vSnap server if you periodically create long-term recovery points (for example, monthly).

Deploying Data Protection for VMware

The offload operation uses data movers that are provided by Data Protection for VMware to backup and restore information to and from the IBM Spectrum Protect server.

For instructions about how to install and configure Data Protection for VMware for integration with IBM Spectrum Protect Plus, see [“Data Protection for VMware deployment checklist”](#) on page 8.

Considerations for integrating Data Protection for VMware with IBM Spectrum Protect Plus

Review the considerations to help to smoothly integrate Data Protection for VMware with IBM Spectrum Protect Plus:

- IBM Spectrum Protect Plus uses the Data Protection for VMware VMCLI command line interface (CLI) to run offload operations. Before you select virtual machines for offloading, review the vmname restrictions for the Backup command. Go to the [online product documentation](#), select the product version, and type backup command in the **Search IBM Knowledge Center** field.
- Review the unsupported characters in virtual machine or datacenter names. Go to the [online product documentation](#), select the product version, and type troubleshooting in the **Search IBM Knowledge Center** field.
- Offloading requires that vCenters and Data Protection for VMware vStorage Backup Servers are registered in IBM Spectrum Protect Plus as a pair, or configured so that all datacenters in a vCenter are registered in IBM Spectrum Protect Plus. The scope for offloaded backups is limited to the vCenter that is configured with the vStorage Backup Server. Backup and restore operations are limited to the same vCenter. Alternatively, virtual machines that are selected for backup should be restricted to the datacenters that configured in IBM Spectrum Protect.
- Data movers must meet the following requirements:
 - The data mover and the IBM Spectrum Protect server must be configured with the same time zone.
 - A single data mover supports one command at a time. If a data mover is performing a backup operation, it cannot perform a restore operation until the backup operation completes.
 - A single data mover can support multiple virtual machines. When selecting multiple virtual machines for offloading, select the virtual machines from a single datacenter per SLA policy. The virtual machines will be distributed among the available datamovers, and will offload in a parallel sequence. If creating multiple offload jobs for virtual machines on the same datacenter, schedule the jobs so that they do not overlap.

- For Linux-based vStorage Backup Servers, the user must have sufficient permissions to run a shell as the tdpvmware user. Typically, root is used as the user when registering a Linux-based vStorage Backup Server.
- To enable back up and offloading of virtual machine templates, you must include the VMENABLETEMPLATEBACKUPS option in the data mover options file. A backup of a templates can be completed only as a Full type backup. For more information about configuring your environment for backing up virtual machine templates, see [Vmenabletemplatebackups](#).
- Instant Disk Restore jobs utilizing offloading are not supported.

Data Protection for VMware deployment checklist

Use a checklist to guide you through the tasks that are required to set up Data Protection for VMware for use with IBM Spectrum Protect Plus.

It is assumed that IBM Spectrum Protect Plus and the IBM Spectrum Protect server are operational and that you have an understanding of these applications.

For information about Data Protection for VMware, see the [online product documentation](#).

Deployment checklist

Check	Action	Description
<input type="checkbox"/>	Download the IBM Spectrum Protect for Virtual Environments: Data Protection for VMware installation package.	Obtain the installation package from the Passport Advantage Online web site. Related documentation Obtaining the Data Protection for VMware installation package
<input type="checkbox"/>	Determine the number of data movers that are required to protect your vSphere environment.	A data mover is a software component of Data Protection for VMware that moves data to and from the IBM Spectrum Protect server. Multiple data movers might be required to protect your vSphere environment. To determine the number of data movers that are required, see technote 2007197 . This technote also includes considerations for using virtual or physical machines for data movers and for data mover locality. Related documentation How IBM Spectrum Protect nodes are used in a virtual environment
<input type="checkbox"/>	Review the hardware and software requirements for the Data Protection for VMware components	To view the hardware and software requirements for the Data Protection for VMware components: 1. Go to the product Support Portal and type all requirements in the search field. 2. Open the <i>IBM Spectrum Protect for Virtual Environments - All Requirements Doc</i> to find links to system requirements by version. Related documentation Planning to install Data Protection for VMware
<input type="checkbox"/>	Contact your IBM Spectrum Protect server administrator to create a policy domain for use with IBM Spectrum Protect Plus	The policy domain specifies the number of days to retain backup and archive copies on the server storage.

Check	Action	Description
<input type="checkbox"/>	Install Data Protection for VMware	<p>To install Data Protection for VMware:</p> <ol style="list-style-type: none"> 1. Start the installation program by using the instructions for your operating system: <ul style="list-style-type: none"> • Windows: Double-click the spinstall.exe file. • Linux: From the root of the installation folder, change directories to CD/Linux/DataProtectionForVMware. <p>From a command line, enter the following command:</p> <pre>./install-Linux.bin</pre> 2. In the installation wizard, select Typical Installation for Windows operating systems or Complete from the Install Set list for Linux operating systems and follow the steps in the wizard to complete the installation. <p>Related documentation</p> <p>Installing the Data Protection for VMware components</p>
<input type="checkbox"/>	Configure Data Protection for VMware	<p>When the installation wizard completes, the Data Protection for VMware vSphere GUI configuration wizard opens to enable you to set up communication with the IBM Spectrum Protect server.</p> <p>Follow the instructions in the wizard to complete the configuration. For additional assistance using the configuration wizard, refer to the GUI help.</p> <p>Tips:</p> <ul style="list-style-type: none"> • During the configuration, you are asked to specify the nodes that are used in your Data Protection for VMware environment. On the Node Definition Options page of the configuration wizard, you can enter a prefix to identify the nodes as common nodes. For example, you could enter ISPP for IBM Spectrum Protect Plus. On all node configuration pages of the wizard, the prefix is automatically added to the node name. <p>You also use the Node Definition Options page to select the policy domain to use when registering new nodes. Use the domain that your IBM Spectrum Protect server administrator created for IBM Spectrum Protect Plus.</p> <ul style="list-style-type: none"> • You define the data mover or data movers for your environment on the Data Mover Nodes page of the configuration wizard. <p>Related documentation</p> <p>Configuring a new installation with the wizard</p>

Check	Action	Description
☐	Install and configure additional data movers on remote systems if required	<p>You can install data movers on remote systems to redistribute the backup workload among multiple systems.</p> <p>To install a data mover on a remote system:</p> <ol style="list-style-type: none"> 1. Start the Data Protection for VMware installation program and complete the following steps for your operating system: <ul style="list-style-type: none"> • Windows: Select Advanced Installation > Install the data mover or mount proxy in the configuration wizard. • Linux: Select Custom from the Install Set list in the configuration wizard. Ensure that Data Protection for VMware data mover is selected. This option is selected by default. 2. When the installation is complete, follow the instructions in Setting up the data mover nodes in a vSphere environment to set up the data movers.
☐	Start IBM Spectrum Protect Plus	<p>To start IBM Spectrum Protect Plus:</p> <ol style="list-style-type: none"> 1. Open a web browser and enter the URL for the virtual machine where IBM Spectrum Protect Plus is deployed. For example: <pre data-bbox="743 926 1469 978">http://hostname</pre> <p>Where hostname is the IP address of the virtual machine.</p> 2. Enter your user name and password. If you are logging on to IBM Spectrum Protect Plus for the first time, the initial user name is admin and the initial password is password. You will be prompted to reset the default password. 3. Click Sign in.
☐	Define the Data Protection for VMware vStorage Backup Server to use for offload operations	<p>To offload data from IBM Spectrum Protect Plus, you must register the vStorage Backup Server and associate it with the vCenter that contains the virtual machines that you want to backup.</p> <p>To specify the vCenter and vStorage Backup Server:</p> <ol style="list-style-type: none"> 1. In the IBM Spectrum Protect Plus navigation menu, click Manage Protection > Hypervisors > VMware > Backup. 2. Click Manage vCenter, and then click Add vCenter. 3. In the vCenter Properties section, complete the fields to add a vCenter. 4. In the IBM Spectrum Protect vStorage Backup Server Settings section, click Link to IBM Spectrum Protect and enter the information for the vStorage Backup Server. <p>Related documentation</p> <p>“Adding a vCenter Server instance” on page 71</p>

Check	Action	Description
<input type="checkbox"/>	Create a service level agreement (SLA) policy that includes offloading to the IBM Spectrum Protect server	<p>SLA policies allow storage and virtualization administrators to create customized templates for the key processes that are involved in the creation and use of backup jobs.</p> <p>To configure an SLA policy that offloads to the IBM Spectrum Protect server:</p> <ol style="list-style-type: none"> 1. In the IBM Spectrum Protect Plus navigation menu, click Manage Protection > Policy Overview, and then click Add SLA Policy. 2. In the New SLA Policy pane, configure the backup retention, frequency, and target site settings and the backup storage replication policy in the Operational Protection section. <p>Tip: If Offload to IBM Spectrum Protect is selected, IBM Spectrum Protect Plus offloads data to the IBM Spectrum Protect server. The retention policy that is defined in the server policy domain created by your IBM Spectrum Protect server administrator is used for offloads to the server rather than the policy that is set in the Retention field. The policy in the Retention field applies only to backups to the vSnap server.</p> <ol style="list-style-type: none"> 3. In the Additional Protection section, click Offload to IBM Spectrum Protect 4. Set the frequency and start time for offloading data to the IBM Spectrum Protect server. 5. Select Leverage most recent backup to enable the offload to occur from the ESX original host or cluster directly. The latest backup image on the vSnap server is mounted. Note that incremental backups are not supported if this option is selected. <p>Related documentation</p> <p>“Creating an SLA policy” on page 67.</p>
<input type="checkbox"/>	Create a backup job definition that uses the SLA policy	<p>To create a backup job definition that uses the SLA policy:</p> <ol style="list-style-type: none"> 1. In the IBM Spectrum Protect Plus navigation menu, click Manage Protection > Hypervisors > VMware > Backup. 2. Select a vCenter and click Select SLA Policy to add the SLA policy to the job definition <p>Related documentation</p> <p>“Backing up VMware data” on page 75.</p>

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 needs.

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](#)

Chapter 2. Installing IBM Spectrum Protect Plus

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

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.

The specifications that are provided in this section can change over time. For the most current systems specifications, 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 Blueprint](#).

Component requirements

Ensure that you have the required system configuration and a supported browser to deploy and run IBM Spectrum Protect Plus. These requirements apply to all installations of IBM Spectrum Protect Plus.

IBM Spectrum Protect Plus support for third-party platforms, applications, services, and hardware parallels that of 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 it at the same level.

Virtual appliance

IBM Spectrum Protect Plus is installed as a virtual appliance.

The IBM Spectrum Protect Plus virtual appliance has five virtual disks that total 370 GB of storage.

Use a Network Time Protocol (NTP) server to synchronize the time zones across resources that are in your environment, such as the IBM Spectrum Protect Plus virtual 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, restore, 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](#).

For initial deployment, the virtual appliance must meet the following minimum requirements:

- 64-bit 8-core machine
- 48 GB memory

Before deploying to the host, ensure that the following requirements are met:

- The correct VMware or Microsoft Hyper-V template.
- vSphere 5.5, 6.0, 6.5, or 6.7 or Microsoft Hyper-V Server 2016 is installed.

For later versions of vSphere, the vSphere Web Client might be required to deploy IBM Spectrum Protect Plus virtual appliances.

- Network information and VMware host information.
- Either an available static IP address to use or access to the Dynamic Host Configuration Protocol (DHCP).

Browser support

Run IBM Spectrum Protect Plus from a computer that has access to the installed virtual appliance. IBM Spectrum Protect Plus was tested against the following web browsers. Later browser versions also might be supported.

- Firefox 55.0.3
- Google Chrome 60.0.3112
- Microsoft Edge 40.15063

If your screen resolution is less than 1024 x 768 pixels, some items might not fit on the window. Pop-up windows must be enabled in your browser to access the help system and some IBM Spectrum Protect Plus operations.

IBM storage requirements

To integrate IBM Spectrum Protect Plus with IBM Spectrum Protect, IBM Spectrum Protect Version 8.1.0 or later is required.

IBM Spectrum Protect Plus ports

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

Port	Protocol	Firewall	Service	Description
22	TCP	Accept	OpenSSH 5.3 (protocol 2.0)	Used for troubleshooting IBM Spectrum Protect Plus.
443	TCP	Accept	A microservice running a reverse-proxy	Main entry point for the client connections (SSL).
5432	TCP	Blocked	PostgreSQL	SQL RDBMS: Supports job management and some security related data and transactions.
5671	TCP, AMQP	Accept	RabbitMQ	Message framework used to manage messages produced and consumed by the VADP proxy and VMware job management workers. Also facilitates job log management.
5672	AMQP	Blocked	RabbitMQ	Message framework used to manage messages produced and consumed within the IBM Spectrum Protect Plus appliance.

Table 1: Incoming firewall connections (IBM Spectrum Protect Plus appliance) (continued)

Port	Protocol	Firewall	Service	Description
8082	TCP	Blocked	Virgo	Modular Java™ application server. Serves core functions for IBM Spectrum Protect Plus including the REST APIs.
8083	TCP	Blocked	Node.js	JavaScript server. Provides higher level APIs to the user interface leveraging the REST APIs running in Virgo.
8090	TCP	Accept	Administrative Console Framework (ACF)	Extensible framework for system administration functions. Supports plugins that perform operations such as system updates and catalog backup/restore.
8092	TCP	Blocked	ACF Plugin EMI	Supports system update, certificate, and license management.
8093	TCP	Blocked	ACF Plugin Catalog Backup and Recovery	Backs up and restores IBM Spectrum Protect Plus catalog data.
8761	TCP	Accept	Discovery Server	Automatically discovers VADP proxies and is used by IBM Spectrum Protect Plus VM backup operations.
27017	TCP	Blocked	MongoDB	Persists configuration related documents for IBM Spectrum Protect Plus.
27018	TCP	Blocked	MongoDB2	Persists recovery metadata documents for IBM Spectrum Protect Plus.

Table 2: Incoming firewall connections (IBM Spectrum Protect Plus appliance - onboard vSnap server)

Port	Protocol	Firewall	Service	Description
111	TCP	Accept	RPC Port Bind	Allows clients to discover ports that Open Network Computing (ONC) clients require to communicate with ONC servers (internal).
2049	TCP	Accept	NFS	Used for NFS data transfer to and from vSnap (internal).
3260	TCP	Accept	iSCSI	Used for iSCSI data transfer to and from vSnap (internal).
20048	TCP	Accept	NFS	Used for NFS data transfer to and from vSnap (internal).

Table 3: Outgoing firewall connections (IBM Spectrum Protect Plus)

Port	Protocol	Service	Description
22	TCP	OpenSSH 5.3 (protocol 2.0)	Used for SSH communications to remote servers running guest applications components.
25	TCP	SMTP	Email service.
389	TCP	LDAP	Active directory services.
443	TCP	VMware ESXi Host	ESXi host port for managing operations.
443	TCP	VMware vCenter	Client connections to vCenter.
636	TCP	LDAP	Active directory services (SSL)

Table 3: Outgoing firewall connections (IBM Spectrum Protect Plus) (continued)

Port	Protocol	Service	Description
902	TCP	VMware NFC service	Network File Copy (NFC) provides a file-type-aware FTP service for vSphere components. ESXi uses NFC for operations such as copying and moving data between datastores by default.
5985	TCP	Windows Remote Management (WinRM)	Hyper-V and guest applications client connections.
8080	TCP	VADP proxy	Virtual machine data protection proxy.
8900	TCP	vSnap	OVA/Installer version of the intelligent storage framework used as a target for data protection operations.

vSnap 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 at the time that the IBM Spectrum Protect Plus virtual appliance is initially deployed. In larger backup enterprise environments, additional vSnap servers might be required.

For initial deployment, ensure that your virtual machine or physical Linux machine meets the following minimum requirements:

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

Adjust data based on backup capacity for more efficient data deduplication. The general rule is 1 GB for every 1 TB of backup data.

Optionally, a solid-state drive (SSD) improves backup and restore performance.

To improve backup performance, configure the pool to use one or more log devices backed by 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 backed by an SSD.

vSnap server virtual machine installation requirements

Before deploying to the host, ensure you have the following:

- The correct VMware or Microsoft Hyper-V template
- vSphere 5.5, 6.0, 6.5. or 6.7 or Microsoft Hyper-V Server 2016
- For later versions of vSphere, the vSphere Web Client might be required to deploy IBM Spectrum Protect Plus appliances.
- Network information and VMware host information.

- Either an available static IP address to use or access to DHCP.

vSnap server physical installation requirements

The following Linux operating systems are supported for physical vSnap installations:

- CentOS Linux7.3.1611 (x86_64)
- CentOS Linux7.4.1708 (x86_64)
- CentOS 7.1804 (7.5) (x86_64)
- Red Hat Enterprise Linux 7.3 (x86_64)
- Red Hat Enterprise Linux 7.4 (x86_64)
- Red Hat Enterprise Linux 7.5 (x86_64)

vSnap server ports

The following ports are used by vSnap servers. Ports that are marked as Accept use a secure connection (HTTPS/SSL).

Table 4: Incoming vSnap firewall connections

Port	Protocol	Firewall	Service	Description
22	TCP	Accept	SSH	Used for troubleshooting vSnap.
111	TCP	Accept	RPC Port Bind	Allows clients to discover ports that ONC clients require to communicate with ONC servers (internal).
2049	TCP	Accept	NFS	Used for NFS data transfer to and from vSnap (internal).
3260	TCP	Accept	iSCSI	Used for iSCSI data transfer to and from vSnap (internal).
20048	TCP	Accept	NFS	Used for NFS data transfer to and from vSnap (internal).

VADP proxy requirements

In IBM Spectrum Protect Plus, running virtual machine backup jobs through VADP can be taxing on system resources. By creating VADP backup job proxies, you enable load sharing and load balancing for your backup jobs. If proxies exist, the entire processing load is shifted off the IBM Spectrum Protect Plus appliance and onto the proxies. Your system must meet the following requirements if you plan to deploy a VADP proxy.

This feature was tested only for SUSE Linux Enterprise Server and Red Hat environments. It is supported only in 64-bit quad core configurations with a minimum kernel of 2.6.32.

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 configurations in the following Linux environments:

- CentOS Linux 6.5+ (beginning with 10.1.1 patch 1)
- CentOS Linux 7.0+ (beginning with 10.1.1 patch 1)
- Red Hat Enterprise Linux 6, Fix pack 4 or later
- Red Hat Enterprise Linux 7, all updates
- SUSE Linux Enterprise Server 12, all updates

The proxy must have the ability to mount NFS file systems, which in many cases requires an NFS client package to be installed. The exact package details vary based on the distribution.

A minimum of 8 GB of RAM is required (16 GB is preferred), along with 60 GB of disk space. Each proxy must have a fully qualified domain name and must be able to resolve and reach vCenter. vSnap servers must be reachable from the proxy. If a firewall is active on the proxy, the following ports on the vSnap server must be reachable (both TCP and UDP): 111, 2049, and 20048.

Port 8080 on the VADP proxy server must be open when the proxy server firewall is enabled. If the port is not open, VADP backups will run on local vmdkbackup instead of the VADP proxy server.

VADP proxy ports

The following ports are used by VADP proxies. Ports that are marked as Accept use a secure connection (HTTPS/SSL).

Table 5: Incoming VADP proxy firewall connections

Port	Protocol	Firewall	Service	Description
22	TCP	Accept	SSH	Port 22 is used to push the VADP proxy to the host node.
8080	TCP	Accept	VADP	VADP REST APIs

Table 6: Outgoing VADP proxy firewall connections

Port	Protocol	Firewall	Service	Description
111	TCP	Accept	vSnap RPC Port Bind	Used for troubleshooting vSnap.
443	TCP	Accept	VMware ESXi Host/vCenter	Allows clients to discover ports that ONC clients require to communicate with ONC servers (internal).

Table 6: Outgoing VADP proxy firewall connections (continued)

Port	Protocol	Firewall	Service	Description
902	TCP	Accept	VMware ESXi Host	NFC provides a file-type-aware FTP service for vSphere components. ESXi uses NFC for operations such as copying and moving data between datastores by default.
2049	TCP	Accept	vSnap NFS	Used for NFS data transfer to and from vSnap (internal).
5671	TCP	Accept	RabbitMQ	Used for iSCSI data transfer to and from vSnap (internal).
8761	TCP	Accept	Discovery Server	Used for NFS data transfer to and from vSnap (internal).
20048	TCP	Accept	vSnap mounted	Mounts vSnap file systems on clients such as the VADP proxy, application servers, and virtualization data stores.

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

VADP proxy on vSnap server requirements

VADP proxies can be installed on vSnap servers in your IBM Spectrum Protect Plus environment. A combination VADP proxy/vSnap server must meet the minimum requirements of both devices. Consult 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 your combination VADP proxy and vSnap server meets the following minimum requirements, which is the sum of the requirements for each device.

VADP proxy installed on a virtual or physical 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.

Hypervisor requirements

Review the hypervisor requirements for IBM Spectrum Protect Plus.

Hyper-V requirements

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

- Microsoft Hyper-V Server 2016
- At least 150 GB of drive space if the disk type is set to fixed size during the installation procedure
- 16 GB memory

Backup and restore of shared virtual hard disks (shared VHDX) is not supported. For known issues and limitations, see [technote 718757](#).

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 started.

Hyper-V servers can be registered using a Domain Name System (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 via DNS. If DNS is not available, the server must be added to the `/etc/hosts` file on the IBM Spectrum Protect Plus virtual appliance by using the command line. 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.

VMware requirements

Ensure that the latest version of VMware Tools is installed in your environment. IBM Spectrum Protect Plus was tested against VMware Tools 9.10.0.

Physical RDM (pRDM) LUNs are not supported for virtual machine backup and restore operations that involve VMware snapshot technology.

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](#).

The following VMware vSphere versions are supported:

- vSphere 5.5 and later
- vSphere 6.0 and later
- vSphere 6.5 and later
- vSphere 6.7 and later

File indexing and restore requirements

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

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.

VMware requirements

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

Windows requirements

IBM Spectrum Protect Plus can protect and restore virtual machines with the following Windows operating systems and file systems. Other file systems are also supported, but only these file systems are eligible for file indexing and restore operations.

Windows operating system and file system requirements

- Windows Server 2008 R2
- Windows Server 2012 R2
- Windows Server 2016

Supported file systems

- NTFS
- ReFS
- CsvFS

File indexing and restore operations support Small Computer System Interface (SCSI) disks in a Hyper-V environment. Integrated Drive Electronics (IDE) disks are not supported. Note that Generation 1 virtual machines require IDE boot disks. However, if additional SCSI disks are available, file indexing and restore operations will be supported on those disks. Windows Remote Shell (WinRM) must be enabled.

When files are indexed in a Windows environment, the following directories on the resource are skipped. Files within these directories are not added to the IBM Spectrum Protect Plus inventory and are not available for file recovery.

- /Drivers
- /Program Files
- /Program Files (x86)
- /Windows
- /winnt

Ensure the latest version of VMware Tools is installed on VMware virtual machines, and Hyper-V Integration Services is installed on Hyper-V virtual machines.

Space requirements

The C drive must have sufficient temporary space to save the file indexing results.

Connectivity requirements

The hostname of the IBM Spectrum Protect Plus appliance should be resolvable from the Windows virtual machine. If the host name of the appliance is not resolvable, add the IP address of the appliance in the `ecxAddress` field in the IBM Spectrum Protect Plus configuration file, which can be found in the following location on the appliance:

```
/opt/virgo/repository/ecx-usr/deploy.vmware.vmdeployer.json
```

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

- The user identity must have the **Log on as a service**. 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 follows the `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 will override the default Windows NTLM protocol. Note that Kerberos

does not allow local user accounts to be used and is only suitable for environments in which all machines are on a single domain.

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.

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

IBM Spectrum Protect Plus can protect and restore virtual machines with the following Linux operating systems and file systems. Other file systems are also supported, but only these file systems are eligible for file indexing and restore operations.

Supported operating systems

- Red Hat Enterprise Linux 6.4+
- CentOS 6.4+
- Red Hat Enterprise Linux 7.0+
- CentOS 7.0+
- SUSE Linux Enterprise Server 12.0+

Supported file systems

- 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.

When files are indexed in a Linux environment, the following directories on the resource are skipped. Files within these directories are not added to the IBM Spectrum Protect Plus inventory and are not available for file recovery.

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

Files in virtual file systems like /proc, /sys, and /dev are also skipped.

Software requirements

Python version 2.6 (any level) or 2.7 (any level) must be installed.

When file systems are indexed, temporary metadata files are generated under the /tmp directory and then deleted as soon as the indexing is complete. The amount of free space required for the metadata depends on the total number of files present on the system. Ensure that there is approximately 350 MB of free space per 1 million files.

For Red Hat Enterprise Linux/CentOS 6.x, ensure that the `util-linux-ng` package is up-to-date by running the `yum update util-linux-ng` command. Depending on your version or distribution, the package may be named `util-linux`.

If data resides on Logical Volume Manager (LVM) volumes, ensure 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` command to update the package if necessary.

If data resides 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 resignature 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
```

For a discussion of the `lvmetad` service, see [The Metadata Daemon \(lvmetad\)](#).

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 Universally Unique Identifier (UUID) is modified. To resolve this issue, update `xfsprogs` to version 4.2.0 or above. 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 SSH. The secure file transfer protocol (SFTP) subsystem for SSH must also be enabled. For SFTP configuration information, see [File Transfer with SFTP](#).

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 set.

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

1. Create user: `useradd -m agent`

Where `agent` specifies the IBM Spectrum Protect Plus Agent user.

2. Set a password: `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
```

Db2 requirements

Review the IBM Db2 database backup and restore requirements for IBM Spectrum Protect Plus.

IBM Db2 configuration requirements

The following Db2 databases are supported:

- IBM Db2 Version 10.5 and later maintenance levels and mod levels: Enterprise Server Edition only.
- IBM Db2 Version 11.1 and later maintenance levels and mod levels: Enterprise Server Edition only.

Operating systems

The following operating systems are supported:

- On PowerPC:
 - AIX 7.1 and later mod and fix pack levels (64-bit kernel)
 - AIX 7.2 and later mod and fix pack levels (64-bit kernel)
- On Linux x86_x64:
 - Red Hat Enterprise Linux 6.8 and later maintenance levels and mod levels: Server Edition
 - Red Hat Enterprise Linux 7 and later maintenance levels and mod levels: Server Edition
 - SUSE Linux Enterprise Server 11.0 SP4 and later maintenance levels and mod levels: Server Edition
 - SUSE Linux Enterprise Server 12.0 SP1 and later maintenance levels and mod levels: Server Edition

Additional notes

To help optimize performance, install the latest IBM Db2 patches and updates in your environment.

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

- Db2 is configured for single partitioned databases only. Backups of multipartition Db2 databases are not supported.
- Db2 archive logging is activated.
- Logical volumes of Db2 table spaces (data and temporary table spaces), local database directory, and Db2 log files are managed by Linux Logical Volume Manager (LVM2) or by AIX Logical Volume Manager (JFS2).
- While Db2 table spaces and the local database directory can reside on the same logical volumes, you must ensure that the Db2 log files reside on separate, dedicated logical volumes (or in a volumes group).
- Db2 pureScale is not supported.

Software

Review the following software requirements:

- The bash and sudo packages must be installed. Sudo must be version 1.7.6p2 or later. Run `sudo -V` to check the version.
- Python version 2.6 (any level) or 2.7 (any level) must be installed.
- Ensure the supported version of Linux x86_64, or AIX® is installed.
- LVM2 on Linux and JFS2 on AIX are used for creating temporary volume snapshots. For more information, see [“LVM2 and JFS2” on page 102](#).
- On AIX, ensure that the NFS communication is configured with reserved ports by using the following command:

```
nfsso -p -o nfs_use_reserved_port=1
```

Connectivity

Review the following connectivity requirements:

- 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 SFTP subsystem for SSH is enabled.

- The server can be registered by using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Protect Plus.

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

The IBM Spectrum Protect Plus agent user must have the following privileges:

- Privileges to run commands as the root user and as the Db2 software owner user by using sudo. IBM Spectrum Protect Plus requires this 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. You can set these privileges in `/usr/local/bin/db2ls` file. IBM Spectrum Protect Plus requires this to discover and collect information about Db2 instances and databases.

Db2 ports

The following ports are used by Db2 servers. Port that are marked as Accept use a secure connection (HTTPS/SSL).

Table 7: Incoming Db2 server firewall connections

Port	Protocol	Firewall	Service	Description
22	TCP	Accept	SSH	Used for SSH data transfer to and from the internal vSnap server.

Table 8: Outgoing Db2 server firewall connections

Port	Protocol	Service	Description
111	TCP	vSnap RPC Port Bind	Allows clients to discover ports that Open Network Connectivity (ONC) clients require to communicate with ONC servers.
2049	TCP	vSnap NFS	Used for NFS file sharing via vSnap.
20048	TCP	vSnap NFS Mount	Mounts vSnap file systems on clients such as the VADP proxy, application servers, and virtualization data stores.

Oracle requirements

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

Oracle configuration requirements

The following Oracle databases are supported:

- Oracle 11g R2
- Oracle 12c R1
- Oracle 12c R2

Operating systems

The following operating systems are supported:

- IBM AIX 6.1 TL9+
- AIX 7.1+
- Red Hat Enterprise Linux / Centos 6.5+
- Red Hat Enterprise Linux / Centos 7.0+
- SUSE Linux Enterprise Server 11.0 SP4+
- SUSE Linux Enterprise Server 12.0 SP1+

Additional notes

Review the following considerations:

- 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 recoveries are not server pool-aware. IBM Spectrum Protect Plus can recover databases to a RAC, but not to specific server pools.

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.
- Python version 2.6 (any level) or 2.7 (any level) must be installed.
- For Red Hat Enterprise Linux/CentOS 6.x, ensure the `util-linux-ng` package is up-to-date by running the `yum update util-linux-ng` command. Depending on your version or distribution, the package might be named `util-linux`.

Connectivity

Review the following 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 using SSH. The secure file transfer protocol (SFTP) subsystem for SSH must also be enabled.
- The server can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Protect Plus.
- When registering Oracle RAC nodes, register each node 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 (referred to as the IBM Spectrum Protect Plus agent user).

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

The IBM Spectrum Protect Plus agent user must have the following privileges:

- Privileges to run commands as a root user and as Oracle software owner users (for example, oracle, grid) using sudo. IBM Spectrum Protect Plus requires these privileges for various tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases and 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. IBM Spectrum Protect Plus requires these privileges to discover and collect 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 user with the required privileges, see [“Sample configuration of an IBM Spectrum Protect Plus agent user” on page 29](#).

NFS requirements

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

During database restore, 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` under each Oracle Home must be owned by a root user and have setuid privileges. This is typically preconfigured by the Oracle installer, but on certain systems, the binary might not have the necessary privileges. For more information, see [Database Startup Failed with Direct NFS](#)

Run the following commands to set the correct privileges:

```
chown root:oinstall ORACLE_HOME/bin/oradism
chmod 750 ORACLE_HOME/bin/oradism
```

Replace oinstall with the appropriate group that owns the installation.

Database discovery

IBM Spectrum Protect Plus discovers Oracle installations and databases by looking through the files `/etc/orainst.loc` and `/etc/oratab`, as well as 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 alternate locations of these 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 such. 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 in order 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 running an associated backup job. If 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 tablespace.

Log backup

- The cron 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 of its own. Privileges can be granted through the `crontab.allow` configuration file.

Sample configuration of an IBM Spectrum Protect Plus agent user

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

- Create the user that will be designated as the IBM Spectrum Protect Plus agent user by using the following command:

```
useradd -m sppagent
```

- Set a password to `passwd sppagent`.
- If you use key-based authentication, place the public key in the `/home/sppagent/.ssh/authorized_keys` file, or in another appropriate file depending on your `sshd` configuration, and ensure the correct ownership and permissions are set, such as:

```
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 by using the following command:

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

- If ASM is in use, also add the user to the OSASM group by using the following command:

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

Oracle ports

The following ports are used by Oracle servers. Note that ports marked as Accept use a secure connection (HTTPS/SSL).

Port	Protocol	Firewall	Service	Description
22	TCP	Accept	SSH	Used for SSH data transfer to and from the internal vSnap server.

Port	Protocol	Service	Description
111	TCP	vSnap RPC Port Bind	Allows clients to discover ports that Open Network Connectivity (ONC) clients require to communicate with ONC servers.
2049	TCP	vSnap NFS	Used for NFS file sharing via vSnap.
20048	TCP	vSnap NFS Mount	Mounts vSnap file systems on clients such as the VADP proxy, application servers, and virtualization data stores.

Microsoft SQL Server requirements

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

Microsoft SQL Server configuration requirements

The following SQL Server database and operating system configurations are supported:

- SQL 2008 R2 SP3 on Windows Server 2012 R2
- SQL 2012 on Windows Server 2012 R2
- SQL 2012 SP2 on Windows Server 2016
- SQL 2014 on Windows Server 2012 R2
- SQL 2014 on Windows Server 2016
- SQL 2016 on Windows Server 2012 R2
- SQL 2016 on Windows Server 2016
- SQL 2017 on Windows Server 2012 R2 (beginning with 10.1.1 patch 1)
- SQL 2017 on Windows Server 2016 (beginning with 10.1.1 patch 1)

For the best performance, install the latest SQL Server patches and updates in your environment.

The following table shows the configurations that are supported by each operating system:

Table 11: Supported configurations by operating system

	Standalone	SQL Server failover clustering	AlwaysOn configurations
SQL 2008 R2 SP3	Yes	Yes	Yes
SQL 2012	Yes	Yes	Yes
SQL 2014	Yes	Yes	Yes
SQL 2016	Yes	Yes	Yes
SQL 2017	Yes (beginning with v10.1.1 patch1)	No	No

Windows Remote Shell (WinRM) must be enabled.

Cluster drive configurations utilizing dynamic disks is not supported. Standalone dynamic disks configurations are supported. 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](#).

In-Memory OLTP requirements and limitations

In-Memory online transaction processing (OLTP) is a memory-optimized database engine used to improve database application performance, supported in SQL 2014 and 2016. The following requirements and limitations apply to in-memory OLTP usage:

- The maximum restore file path must be fewer than 256 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 VSS and SQL restore capabilities.

SQL incremental backup requirements and limitations

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

The following requirements enable write range tracking:

- Windows 2012 R2 or later
- NTFS version 3.0 or later

Cluster shared volumes (CSVFS) are supported.

The following technologies are not supported for changed bytes tracking:

- 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.

The minimum space required for shadow copy storage is 100 MB, though more space may be required on systems with increased activity. The SQL agent checks the source volume space, and will fail a backup if the free space on the source volume is less than 100 MB. A warning message displays in the job log when free space is less than 10%; then the backup proceeds.

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

- Journal discontinuity is detected, due to the log reaching the maximum size, disabling of the log, or changing of the cataloged USN ID.
- The file size is less than the tracked threshold size, which by default is 102 KB.
- The file size is smaller after a backup job.
- A file is added after a backup job.

SQL Always On availability groups considerations

To use Microsoft Server SQL Always On availability groups, configure the preferred instance for backup by using SQL Server Management Studio. Select the Availability Group node. Select the availability group that you want to configure, and then select **Properties**. In the **Availability Group Properties** dialog box, select **Backup Preferences**.

On the **Where should backups occur?** pane, select **Secondary only**. If more than one secondary replica is available, the SQL Executor will select the first secondary replica in the preferred list reported by the IBM Spectrum Protect Plus SQL agent. The SQL agent sets the VSS backup type to COPY_ONLY.

Registration and authentication

Register each SQL Server server in IBM Spectrum Protect Plus by name or IP address. When registering a SQL Server Cluster (Always On) node, register each node by name or IP address. 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 routeable 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 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.

Kerberos requirements

You can enable Kerberos-based authentication by editing 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 user name must be able to authenticate with 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.

SQL log truncation

To truncate Microsoft SQL Servr application logs during hypervisor backup jobs, enable the **Truncate SQL logs** option when creating a job definition. The following SQL environments support SQL log truncation:

- SQL 2008 R2 SP3 on Windows Server 2012 R2
- SQL 2012 SP4 on Windows Server 2012 R2
- SQL 2014 SP2 on Windows Server 2012 R2
- SQL 2016 SP1 on Windows Server 2016

Privileges

On the SQL Server server, the system login credential must have public and sysadmin permissions enabled, in addition to permission to access cluster resources in a SQL Server Always On 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.

The SQL VDI-based framework is used to interact with SQL databases and log backup/restore operations. A VDI connection requires SQL 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 [SQL Server VDI backup and restore operations require Sysadmin privileges](#).

The target SQL server service account must have permissions to access SQL restore files. See the "Administrative Considerations" section in [Securing Data and Log Files](#).

To back up logs, the SQL Server user registered with IBM Spectrum Protect Plus must have the sysadmin permission enabled to manage SQL Server agent jobs.

SQL Server ports

The following ports are used by Microsoft SQL Server servers. The ports that are marked as Accept use a secure connection (HTTPS/SSL).

Table 12: Incoming SQL Server firewall connections

Port	Protocol	Firewall	Service	Description
5985	TCP	Accept	WinRM	Windows Remote Management Service
5986	TCP	Accept	WinRM	Secure Windows Remote Management Service

Table 13: Outgoing SQL Server firewall connections

Port	Protocol	Service	Description
3260 iSCSI initiator is required on this node.	TCP	vSnap iSCSI	iSCSI vSnap target port used for mounting LUNS for backup and recovery.

Obtaining the IBM Spectrum Protect Plus installation package

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

Before you begin

For the list of installation packages by component, and the links to the download site for the files, see [technote 4044949](#).

Procedure

Download the appropriate installation files.

Different installation files are provided for installation on VMware and Microsoft Hyper-V systems and for installation on physical or virtual machines. Ensure that you download the correct files for your environment.

Related tasks

[“Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 34](#)

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 ESX or ESXi server.

[“Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance” on page 36](#)

To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import a 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 machine.

[“Installing vSnap servers” on page 39](#)

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 ESX or ESXi server.

Before you begin

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in [“Component requirements ” on page 13](#) and [“Hypervisor requirements ” on page 21](#).
- Download the virtual appliance .ova template installation file from Passport Advantage Online. For information about downloading files, see [technote 4044949](#).
- Run MD5 Checksum on the downloaded .ova 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 37](#).

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 by taking one of the following actions:
 - a) If you are using the vSphere Client, from the **File** menu, click **Deploy OVF Template**.
 - b) If you are using the vSphere Web Client, click **Create/Register VM**, then select **Deploy a virtual machine from an OFV or OVA file**
2. Specify the location of the .ova template installation file and select it. Click **Next**.
3. Review the template details and accept the End User License Agreement. 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. Identify the datacenter, server, and resource pool for deployment. When prompted to select storage, select from the datastores that are already configured on the destination host. The virtual machine configuration file and virtual disk files are stored on the datastore. Select a datastore that is large enough to accommodate the virtual machine and all of its virtual disk files. Click **Next**.
6. Select a disk format to store the virtual disks. To help optimize performance, keep the default option, thick provisioning. Thin provisioning requires less disk space, but might impact performance. Click **Next**.
7. Select networks for the deployed template to use. Several available networks on the ESX server might 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**.
8. Enter network properties for the virtual machine default gateway, DNS, IP address, network prefix, and machine host name. You can enter a static IP configuration, or leave all fields blank to use a DHCP configuration. Work with your network administrator when configuring network properties. The network prefix should be specified by a network administrator. The network prefix must be entered using Classless Inter-Domain Routing (CIDR) notation; valid values are 1 - 32.
9. Click **Next**.
10. Review your template selections. 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: The virtual machine must remain powered on for the IBM Spectrum Protect Plus application to be accessible.
12. Record the IP address of the newly created VM.

The IP address is required to log on to the application. Find the IP address in vSphere Client by clicking your newly created VM and looking in the **Summary** tab.

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

What to do next

To run IBM Spectrum Protect Plus, access the newly created VM. A local vSnap server that is already named and registered is also installed on the VM. To change the IP address allocation type after IBM Spectrum Protect Plus deploys, redeploy the virtual machine. After you install the virtual appliance, complete the following actions:

Action	How to
If you use a static IP address instead of DHCP, restart the virtual appliance.	Refer to the documentation for the virtual appliance.
Upload the product key.	See “Uploading the product key” on page 38 .
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 52 .

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

To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import a 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 machine.

Before you begin

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in [“Component requirements”](#) on page 13 and [“Hypervisor requirements”](#) on page 21.
- Download the .exe installation file from Passport Advantage Online. For information about downloading files, see [technote 4044949](#).
- Review additional Hyper-V system requirements. See [System requirements for Hyper-V on Windows Server 2016](#).
- Run MD5 Checksum on 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.
- 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.
- 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 server starts.

Procedure

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

1. Copy the .exe installation file to your Hyper-V server.
2. Start the installer, complete the installation steps, and then close the installer.
3. Open Hyper-V Manager and select the required server.
4. From the **Actions** menu in Hyper-V Manager, click **Import Virtual Machine**, and then click **Next**. The Locate Folder dialog opens.
5. Browse to the location that you designated during the installation and select the Virtual Machines folder.
6. Click **Next**. The Select Virtual Machine dialog opens.
7. Select **SPP-{release}**, 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, click **Hard Drive**.
13. Click **Edit**, then click **Next**.
14. In the Choose Action screen, click **Convert** and then click **Next**.
15. For the Disk Format, click **VHDX**.
16. For the Disk Type, click **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. After the conversion is completed, click **Browse**, and then select the newly created VHDX.

20. Repeat steps 15 through 20 for each disk under the SCSI Controller section.
21. Power on the virtual machine from Hyper-V Manager.
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 nmtui tool.
For more information, see [“Assigning a static IP address” on page 37](#).

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 38 .
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 52 .

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.
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 34](#)

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 ESX or ESXi server.

[“Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance” on page 36](#)

To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import a 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 machine.

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`.

You are prompted to enter a new password to access the Administrative Console upon first log in.

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 52 .

Chapter 3. Installing and configuring 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. The default 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 default 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 Blueprint](#).

The 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. Register the vSnap server as a backup storage target in IBM Spectrum Protect Plus.
3. Initialize the system and create a storage pool.

Installing vSnap servers

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 13.
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 4044949](#).

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 `.run` installation file from Passport Advantage Online. For information about downloading files, see [technote 4044949](#).

5. Before running the vSnap installation file, ensure that your system is up to date by running the yum update command.
6. Make the file executable through the command `chmod +x file_name.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 “Managing vSnap servers” on page 42.

Installing a virtual vSnap server in a VMware environment

To install a virtual vSnap server in a VMware environment, deploy an OVF template. This creates a virtual appliance containing the vSnap server on a VMware virtual machine.

Procedure

1. Download the vSnap server .ova template installation file from Passport Advantage Online. For information about downloading files, see [. .](#)
2. To deploy the vSnap server, take one of the following actions:
 - If using the vSphere Client to deploy the vSnap server, from the **File** menu, choose **Deploy OVF Template**.
 - If using the vSphere Web Client, click **Create/Register VM**, then select **Deploy a virtual machine from an OVF or OVA file**. Click **Next**.
3. Specify the location of the IBM Spectrum Protect Plus OVA template installation file and select it. Click **Next**.
4. Review the template details and accept the End User License Agreement. Click **Next**.
5. 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**.
6. Identify the datacenter, server, and resource pool for deployment. When prompted to select storage, select from the datastores that are already configured on the destination host. The virtual machine configuration file and virtual disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine and all of its virtual disk files. Click **Next**.
7. 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**.
8. 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**.
9. Enter network properties for the virtual machine default gateway, DNS, IP address, network prefix, and machine host name. You can enter a static IP configuration, or leave all fields blank to use a DHCP configuration. Work with your network administrator when configuring network properties.

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 - 32.
10. Click **Next**.
11. Review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template. Deployment might take significant time.
12. 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.

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

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

What to do next

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

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See “Managing vSnap servers” on page 42.

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 .exe vSnap installation file from Passport Advantage Online. For information about downloading files, see [technote 4044949](#).
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 2016](#).
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 37](#).

What to do next

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

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See “Managing vSnap servers” on page 42 .

When uninstalling IBM Spectrum Protect Plus in a Hyper-V environment, it is recommended to delete the IBM Spectrum Protect Plus appliance from Hyper-V first before running the uninstaller.

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 46](#).

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`.
You are prompted to change this password during the first logon.
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 42.

Password

Enter the user name for the vSnap server that you created in step “2” on page 42.

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 44.
Expand the vSnap storage pool.	See “ Expanding a vSnap storage pool ” on page 45.
If necessary, configure and administer certain aspects of vSnap, such as network configuration or storage pool management.	“ vSnap server administration reference ” on page 46

Related tasks

“[Start IBM Spectrum Protect Plus](#)” on page 52

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.

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.

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 47](#).

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 default 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`.

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 46.

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 47.

Setting vSnap storage options

You can set additional storage-related options for a vSnap server.

Procedure

To set the options for a vSnap server, 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, and then expand the **Storage Options** section. Set the storage options.

Enable Compression

If enabled, each incoming block of data is compressed using a compression algorithm before it is written to the storage pool. Compression consumes a moderate amount of additional CPU resources.

Enable Deduplication

If enabled, 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 disabled 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.

Synchronous Write Mode

Disabling synchronous writes can lead to data loss and silent corruption of backup data if the storage server experiences an abrupt shutdown or reboot during a backup job. Do not disable this option unless the storage server resides in a stable environment that is adequately secured against hardware and power failures.

Encryption Enabled

This option displays the encryption status of the vSnap server. Encryption can be enabled only during vSnap initialization. This option is for informational purposes only.

3. Click **Save**.

Expanding a vSnap storage pool

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.

Procedure

To set the 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.

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 67

vSnap server administration reference

Once the vSnap server has been 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.

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

Managing vSnap 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`.

By default, the `serveradmin` user is not assigned sudo privileges. To assign sudo privileges to the `serveradmin` user, log in to the vSnap server command line interface and enter the following command:

```
echo "serveradmin ALL=(ALL) NOPASSWD: ALL" >/etc/sudoers.d/serveradmin
```

The command line interface consists of several commands and subcommands that manage various aspects of the system. Refer to [Storage Management on page 57](#) and [Network Management on page 59](#) 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 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 manage it.

Related tasks

[“Installing vSnap servers” on page 39](#)

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.

[“Managing vSnap servers” on page 42](#)

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.

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.

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 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

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 44](#), 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 48](#). 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 48.

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 (currently used for VMware backups).

iscsi

This service is used for data traffic when backing up data using iSCSI (currently used for Hyper-V backups).

smb

This service is used for data traffic when backing up data using SMB/CIFS (currently not used, reserved for future use.)

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.

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 > Backup**.

Procedure

1. Log on to the vSnap server console with the user ID `serveradmin`. The initial password is `sppDP758`.
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 46.
2. Run the following commands:

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

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.

Chapter 4. Getting off to a quick start

To start using IBM Spectrum Protect Plus, you must complete steps that include defining the location of the resources that you want to protect and creating backup policies for those resources. This getting started section provides the basic steps to set up and start using IBM Spectrum Protect Plus.

Before you start, ensure that the tasks that are listed in [“Product deployment roadmap” on page 4](#) are complete.

In this section, 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.

Then, hypervisor and database application backup and restore tasks are completed by the *application administrator* as shown in the following table. However, a single administrator might be responsible for all tasks in your environment.

Action	Owner	
Start IBM Spectrum Protect Plus	Infrastructure administrator and application administrator	The infrastructure administrator starts the application for the first time by using the default <code>admin</code> user account. The application administrator starts the application by using the user account that is created by the infrastructure administrator.
Create backup policies	Infrastructure administrator	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 offload backup data to secondary backup storage for longer-term protection. Backup policies are called service level agreement (SLA) policies in IBM Spectrum Protect Plus.
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” on page 56	Application administrator	Resources are servers for hypervisors or database applications that host data that you want to protect.

Action	Owner	
“Add resources to a job definition” on page 58	Application administrator	Job definitions associate the resources that you want to back up with one or more SLA policies. The options and schedules that are defined in the SLA policies are used for backup jobs for the resources.
Start a job	Application administrator	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 initial user name is `admin` and the initial password is `password`. You will be prompted to reset the default password.
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`. The `serveradmin` user is used to access the administrative console and the IBM Spectrum Protect Plus virtual appliance. The the password for `serveradmin` must be changed before accessing the administrative console and IBM Spectrum Protect Plus virtual appliance.
 - Start the initialization process for the onboard vSnap server. Select **Initialize** or **Inititalize with encryption enabled** to encrypt data on the server.

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 offloading to IBM Spectrum Protect, which are optional features. For information about how to set up these features in the SLA policy, see [“Creating an SLA policy” on page 67](#).

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 or the number of snapshots to keep.

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.

If multiple sites are available to IBM Spectrum Protect Plus, the vSnap backup destination with the largest amount of available storage is used first.

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:

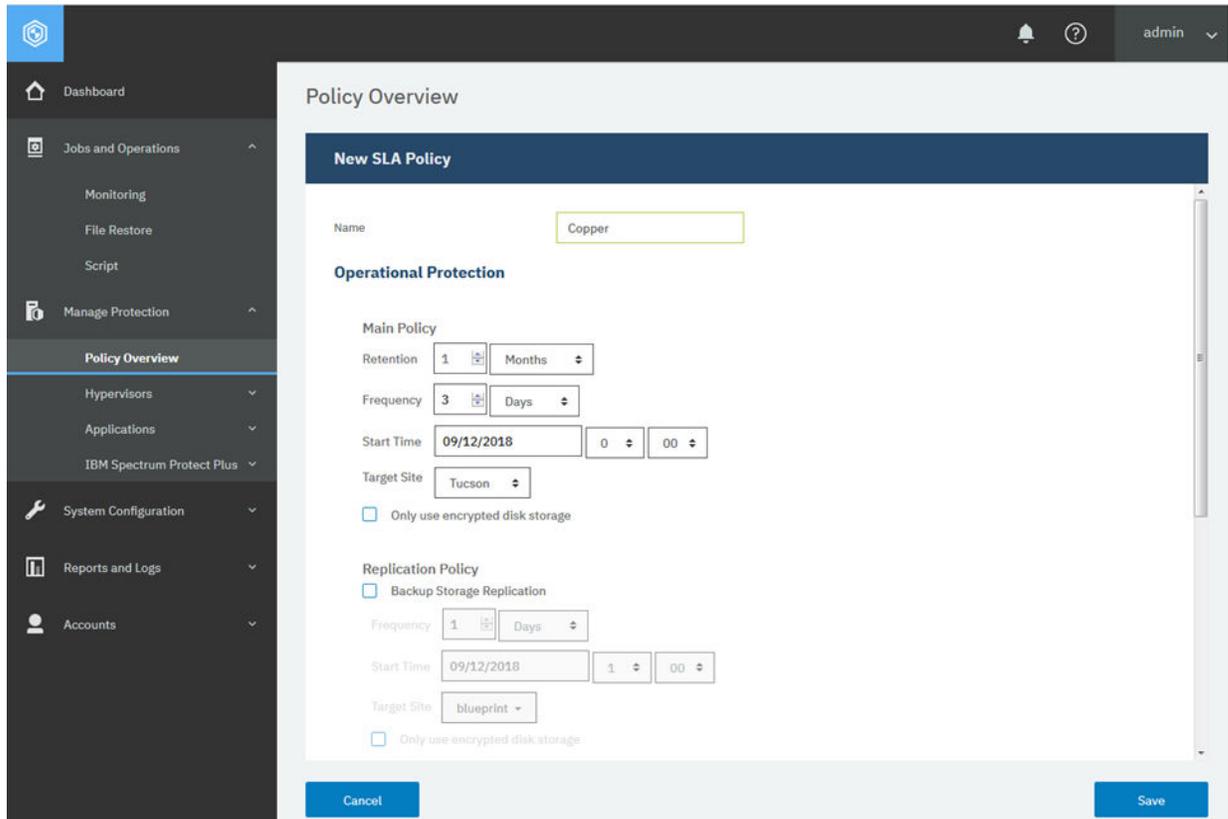


Figure 2: Creating an SLA policy

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 58.

Related concepts

[“Replicate backup-storage data ”](#) on page 6

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.

[“Offload to IBM Spectrum Protect ”](#) on page 6

Although the vSnap server is the primary backup location for snapshots, you can also offload snapshots to IBM Spectrum Protect.

[“Managing SLA policies for backup operations”](#) on page 67

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 and the option to offload backup data.

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 177.

To create custom user roles and resource groups, see [“Creating a resource group”](#) on page 170 and [“Creating a role”](#) on page 174

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.

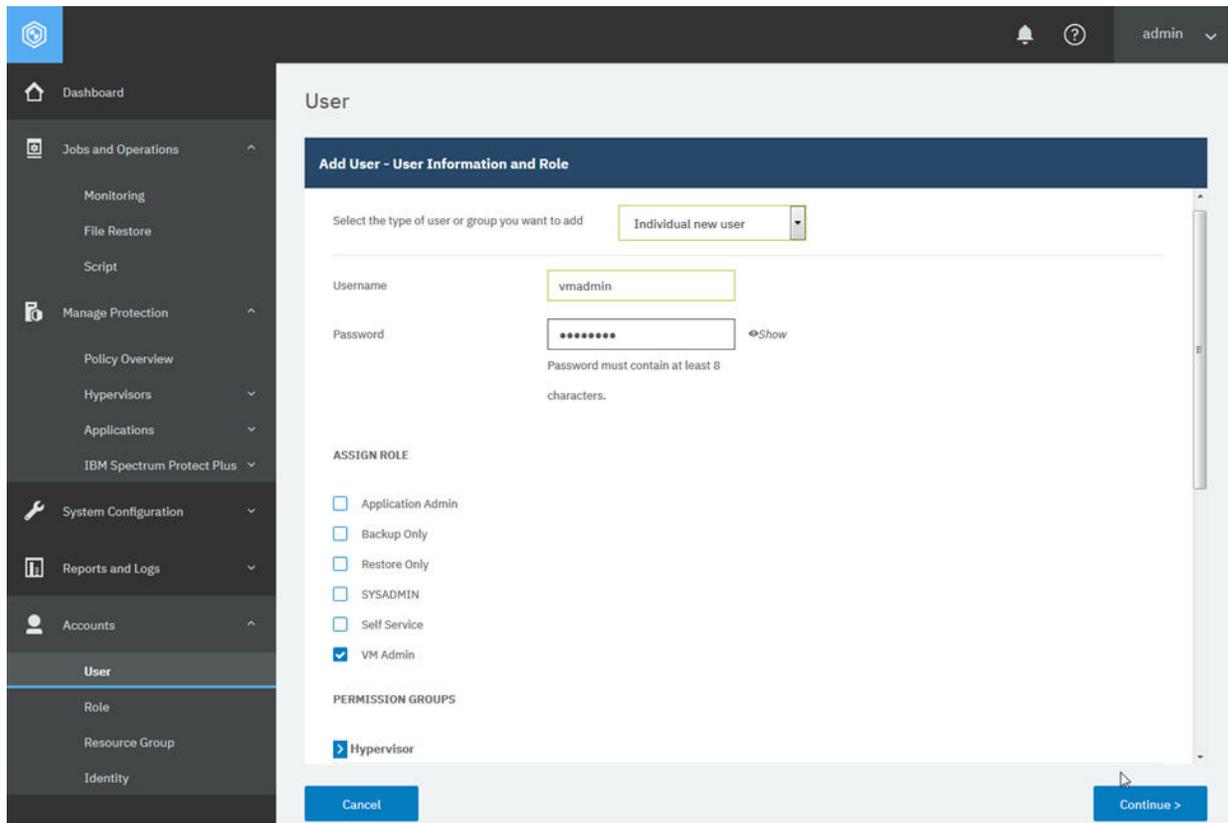


Figure 3: Creating a user account and assigning a role

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.

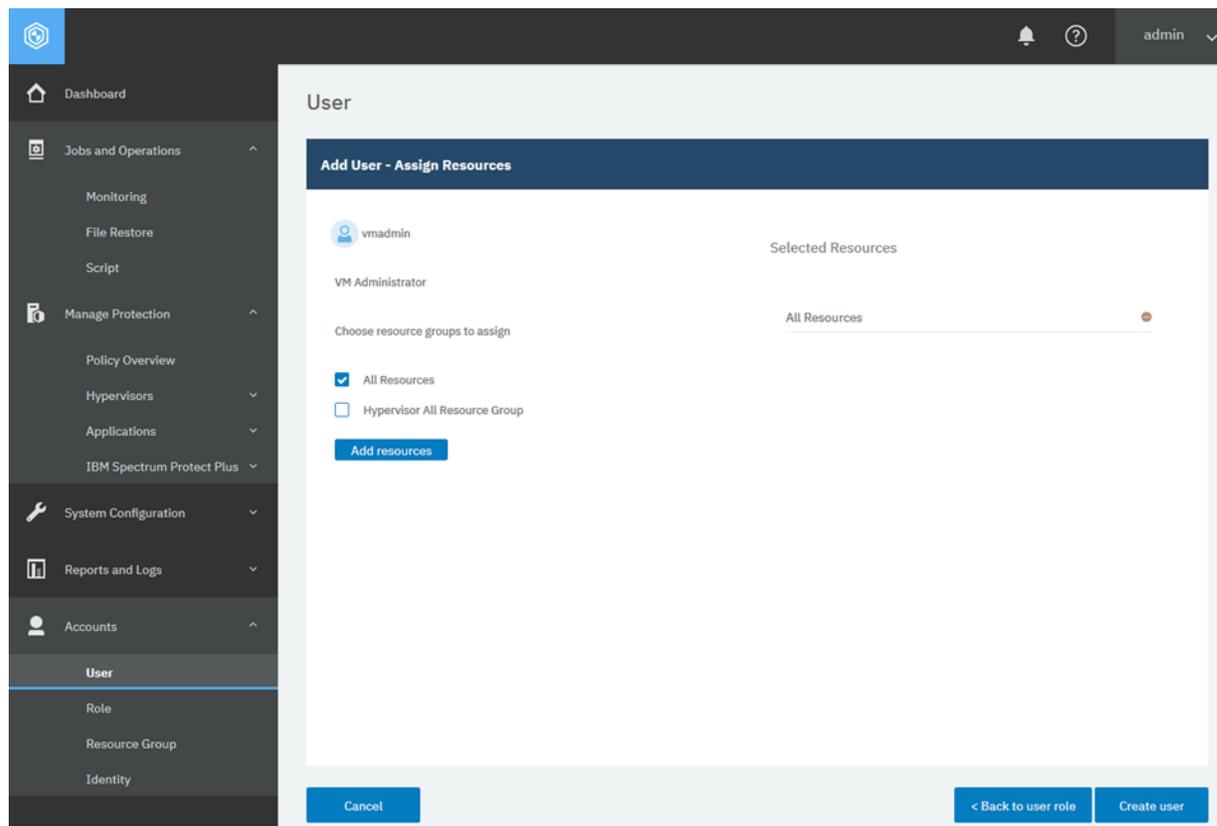


Figure 4: Selecting a resource group for the user account

8. Click **Create user**.

Related concepts

“Managing user access” on page 169

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 7, “Protecting hypervisors,” on page 71 and Chapter 8, “Protecting applications,” on page 101.

Procedure

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

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

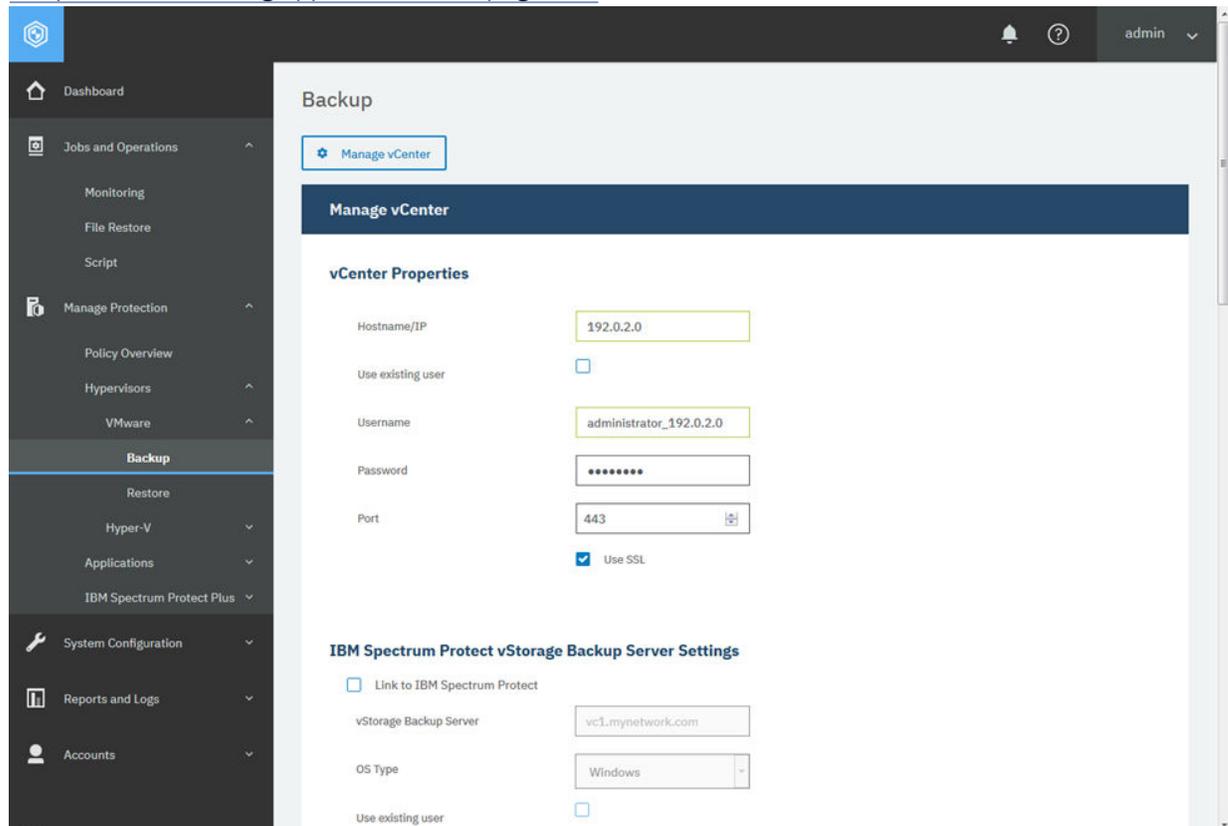
Password

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.

The following example shows populated fields. Note that the **Link to IBM Spectrum Protect** option is not selected in this example. This option enables you to configure an associated vStorage Backup Server for IBM Spectrum Protect offload functionality. For more information about this option, see the add resource instructions by resource type in [Chapter 7, “Protecting hypervisors,” on page 71](#) and [Chapter 8, “Protecting applications,” on page 101](#).



The screenshot displays the 'Manage vCenter' configuration page in the IBM Spectrum Protect Plus web interface. The left sidebar shows the navigation menu with 'Backup' selected. The main content area is titled 'Backup' and contains a 'Manage vCenter' button. Below this, the 'vCenter Properties' section includes the following fields: Hostname/IP (192.0.2.0), Use existing user (unchecked), Username (administrator_192.0.2.0), Password (masked with dots), and Port (443). The 'Use SSL' checkbox is checked. The 'IBM Spectrum Protect vStorage Backup Server Settings' section includes the 'Link to IBM Spectrum Protect' checkbox (unchecked), vStorage Backup Server (vc1.mynetwork.com), OS Type (Windows), and another 'Use existing user' checkbox (unchecked).

Figure 5: Adding a vCenter Server instance

- Optional: To configure additional options, including options for offloading data to IBM Spectrum Protect, see [“Adding a vCenter Server instance” on page 71](#).
- 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 7, “Protecting hypervisors,”](#) on page 71 and [Chapter 8, “Protecting applications,”](#) on page 101.

Procedure

To select an SLA policy, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware > Backup**.
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**, and **Tags & Categories**. Tags, which are applied in vSphere, make it possible assign metadata to virtual machines.

The following example shows a hard disk in virtual machine CETVM01WIN selected for backup:

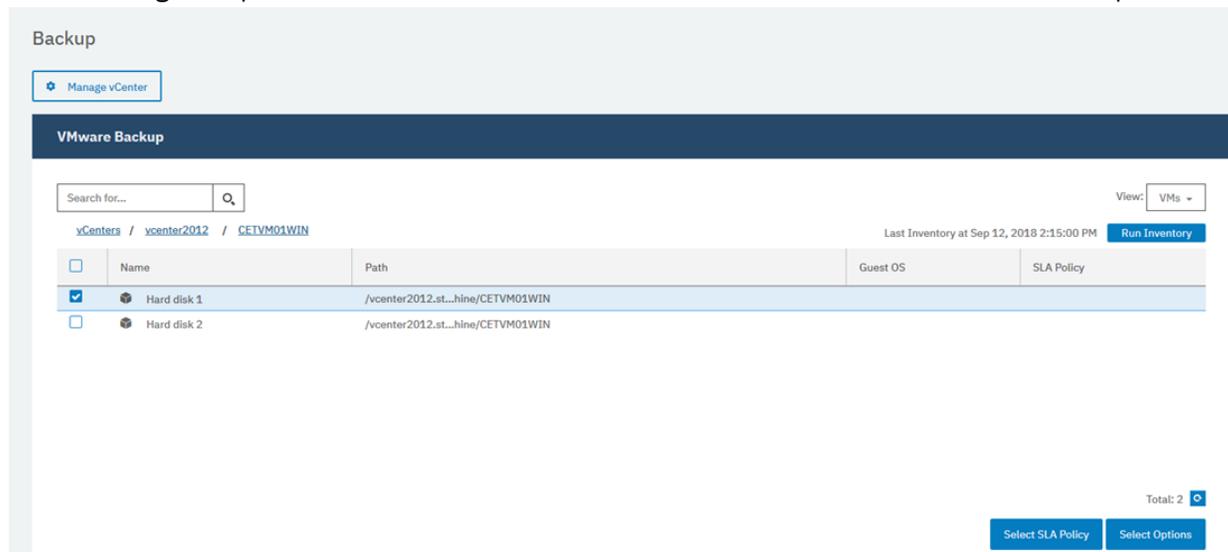


Figure 6: Selecting resources for backup

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:

SLA Policy

	SLA Policy	Frequency	Retention
<input type="checkbox"/>	Gold	Every 4 Hours	1 Weeks
<input type="checkbox"/>	Silver	Every 1 Days at 10:10:00 PM	1 Months
<input type="checkbox"/>	Bronze	Every 1 Days at 10:10:00 PM	1 Weeks
<input checked="" type="checkbox"/>	Copper	Every 3 Days at 12:00:00 AM	1 Months

Save

Figure 7: Selecting an SLA policy

4. To create the job definition by using default options, click **Save**.
5. Optional: To configure additional options, click **Select Options** and follow the instructions in [“Backing up VMware data”](#) on page 75.
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 79.

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 > Monitoring** and starting the job from the **Job Monitor** pane. For instructions, see [“Start a job”](#) on page 59.

Related concepts

[“Protecting IBM Spectrum Protect Plus”](#) on page 141

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, search data, and job information are backed up to a vSnap server defined in the associated SLA policy.

Start a job

You can start a job on demand outside of the schedule that is set by the SLA policy.

Procedure

To start a job on demand, complete the following steps:

1. From the navigation menu, click **Jobs and Operations > Monitoring**.
2. To start the job session, click the **Actions** menu that is associated with the job that you want to start and click **Start** as shown in the following example:

Job Monitor

Filter Job Log: INFO x WARN x ERROR x

Name	Type	Status	Next Run	Last Run	Actions
vmware_Copper	Backup	IDLE	Sep 15, 2018 12:00:00 AM	Sep 12, 2018 12:00:00 AM	<ul style="list-style-type: none"> Start Pause Schedule

Figure 8: Starting a job

What to do next

Optionally, to view job session details such as the job start time, end time, and status, expand the running job session.

Related concepts

[“Managing data protection jobs” on page 143](#)

You can run jobs on demand, pause or cancel running jobs, and pause scheduled 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:

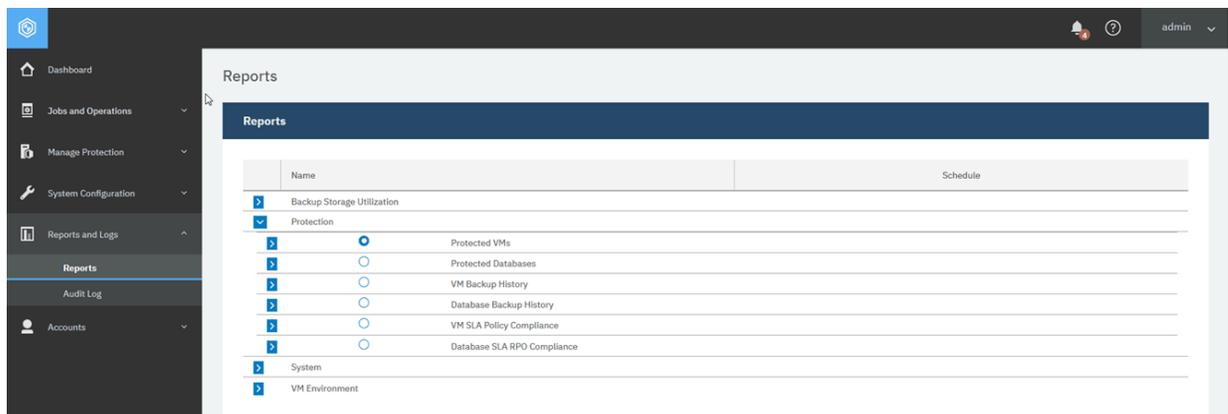


Figure 9: Selecting a report to run

3. Run the report either with custom parameters or default parameters:
 - To run the report with custom parameters, click **Options**, set the report parameters, 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 161](#)

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 5. Updating IBM Spectrum Protect Plus components

You can update the IBM Spectrum Protect Plus virtual appliance, the vSnap server, and the VADP proxy 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 of a component.

For information about available update files and how to obtain them from an IBM download site, see [technote 4044949](#).

Review the following restrictions and tips:

- You must separately update vSnap servers or VADP proxies that are not on IBM Spectrum Protect Plus virtual appliances.
- A patch might not require updates for all IBM Spectrum Protect Plus components. The update files that are available in each patch might vary.
- 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 141.
- 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.
- Ensure no active jobs are scheduled to run during the IBM Spectrum Protect Plus virtual appliance update procedure. After associated jobs are completed or are in an idle state, navigate to **Jobs and Operations > Monitoring**, and then select **Pause Schedule** from the **Actions** list for each job. After the update is completed, select **Release Schedule** for each job to resume the associated schedule.

Updating the IBM Spectrum Protect Plus virtual appliance

Use the IBM Spectrum Protect Plus administrative console to update the virtual appliance.

Before you begin

Updating your IBM Spectrum Protect Plus environment from version 10.1.1 to version 10.1.2 requires a prerequisite update of the version 10.1.1 IBM Spectrum Protect Plus application as well as a full update of the IBM Spectrum Protect Plus virtual appliance and operating system.

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 141.

2. Download the prerequisite IBM Spectrum Protect Plus update named **CNUQ6EN.iso** to a directory on the computer that is running the browser for the administrative console. This .ISO update file will be installed first.
3. Download the full update of the IBM Spectrum Protect Plus virtual appliance and operating system named **CNUQ5EN.iso** to a directory on the computer that is running the browser for the administrative console. This .ISO update file will be installed second.

For a list of download images, including the required operating system update for the virtual appliance, see [technote 4044949](#).

Procedure

To update the IBM Spectrum Protect Plus virtual appliance, complete the following steps:

1. From a supported web browser, access the administrative console at 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.
System (recommended)	To login as a system user, enter the system password. The default password is sppadLG235. You are prompted to change this password during the first login.

3. Click **Manage updates**.
4. Click **Browse** to browse for the prerequisite update file named **CNUQ6EN.iso**, and then click **Upload Update Image**.
The update process begins once the update image is uploaded to the appliance.
5. When the update completes, the virtual machine where the application is deployed automatically restarts. Clear your browser's cache before applying the next update.
6. Start the administrative console, and click **Manage updates**.
7. Click **Browse** to browse for the full update of the IBM Spectrum Protect Plus virtual appliance and operating system named **CNUQ5EN.iso**, and then click **Upload Update Image**.
The update process begins when the update image is uploaded to the appliance.
8. When the update completes, the virtual machine where the application is deployed automatically restarts.



Attention: If upgrading from IBM Spectrum Protect Plus version 10.1.1 and not version 10.1.1p1, you must manually reboot the virtual appliance. If the virtual appliance is not rebooted manually, the upgrade will fail.

HTML content from previous versions of IBM Spectrum Protect Plus might be stored in your browser cache. Clear the cache before logging in to an updated version of IBM Spectrum Protect Plus to ensure that you are viewing the latest content changes.

After updating to V10.1.2, you will be prompted to change the `serveradmin` password upon logging in to the IBM Spectrum Protect Plus user interface. The `serveradmin` user is used to access the Administrative Console and the IBM Spectrum Protect Plus appliance. The root user account is no longer valid after updating to V10.1.2. The `serveradmin` password must be changed before accessing the Administrative Console and the IBM Spectrum Protect Plus appliance. The default password is sppDP758.

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

Related tasks

[“Updating vSnap servers” on page 63](#)

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 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

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 141](#).
2. Download the IBM Spectrum Protect Plus version 10.1.2 virtual appliance plus operating system update file to the **tmp** directory on the vSnap server. For information about downloading files, see [technote 4044949](#).
3. Download the vSnap update script and copy it to a temporary location on the vSnap server. For information about downloading files, see [technote 4044949](#).

Procedure

To update a vSnap server, complete the following steps:

1. Ensure that there are no active jobs that use the vSnap server.
2. Navigate to the **Jobs and Operations > Monitoring** page in IBM Spectrum Protect Plus.
3. When the associated jobs are complete or in an idle state, select **Pause Schedule** from the **Actions** list for each job.
4. On the vSnap server, open a terminal.
5. From the folder where the vSnap update script is located, make the script executable by using the following command:

```
chmod 777 script_name.sh
```

6. Execute the vSnap update script by using the following command:

```
./script_name.sh /tmp/full_update.iso
```

7. Reboot the vSnap server when prompted.
After reboot, additional components will be installed which may take several minutes to complete.
8. After the vSnap server update completes, you will be prompted to change the **serveradmin** password. The serveradmin user is used to access the vSnap server command line interface.
The serveradmin user will not have sudo privileges. To assign sudo privileges to the serveradmin user, log in to the vSnap server command line and enter the following command:

```
echo "serveradmin ALL=(ALL) NOPASSWD: ALL" >/etc/sudoers.d/serveradmin
```

9. On the **Monitoring** page, select **Release Schedule** from the **Actions** list for the jobs that are associated with the vSnap server.

Updating vSnap servers in a Red Hat Enterprise Linux environment

When you update the IBM Spectrum Protect Plus appliance, the default vSnap server is updated automatically. However, if additional vSnap servers are installed on either virtual or physical appliances, you must update those vSnap servers separately. If the vSnap server is running Red Hat Enterprise Linux, the operating system must be updated manually before you update the vSnap server software.

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 141.](#)
2. Update the Red Hat Enterprise Linux operating system on the vSnap server to version 7.4 or 7.5. Follow the update procedure in the Red Hat Enterprise Linux environment documentation.
3. Download the vSnap server update file for Linux to a temporary location on the vSnap server. The file has a .run extension. For instructions about downloading update files, see [technote 4044949.](#)

Procedure

To update a vSnap server, complete the following steps:

1. Ensure no active jobs are using the vSnap server by taking the following actions:
 - a. In the IBM Spectrum Protect Plus navigation pane, expand **Jobs and Operations** and click **Monitoring**.
 - b. When the associated jobs are complete or in an idle state, select **Pause Schedule** from the **Actions** list for each job.
2. On the vSnap server, open a terminal.
3. From the folder where the .run file is located, make the file executable and run the installer by using the following commands:

```
chmod +x file_name.run
```

```
./file_name.run
```

The vSnap packages are installed.

4. After the vSnap server update is complete, you are prompted to change the password for the `serveradmin` user. The `serveradmin` user is used to access the vSnap server command-line interface.
5. Change the password for the `serveradmin` user.
6. To resume scheduled jobs, on the **Monitoring** page, select **Release Schedule** from the **Actions** list for the jobs that are associated with the vSnap server.

Updating VADP proxies

The default VADP functionality is updated with the IBM Spectrum Protect Plus appliance. You must update additional VADP proxies separately.

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 141.](#)

Procedure

If a VADP proxy update is available for external proxies during a reboot 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.

What to do next

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

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 75. The proxies are indicated in the job log by a log message similar to the following text: Run <code>remote vmdkbackup</code> of MicroService: <code>http://<proxy</code> <code>nodename, IP:proxy_IP_address</code>
Uninstall the proxies when you cease running the VMware backup jobs.	To uninstall a proxy, run the following command on the host system from the <code>uninstall</code> subdirectory of the installation directory <code>/opt/IBM/SPP</code> : <code>./uninstall_vmdkbackup</code>

Related tasks

[“Creating VADP backup proxies”](#) on page 80

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Chapter 6. Managing SLA policies for backup operations

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 and the option to offload backup data.

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.

Creating an SLA policy

You can create custom SLA policies to define backup frequency, retention, replication, and offload 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.

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 or the number of snapshots to keep.

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 a primary or secondary backup destination.

If more than one primary or secondary storage site is available to IBM Spectrum Protect Plus, the vSnap backup destination with the largest amount of available storage is used first.

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 46.

Backup Storage Replication

Select this option to enable replication.

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 a primary or secondary replication destination.

If more than one primary or secondary storage site is available to IBM Spectrum Protect Plus, the replication destination with the largest amount of available storage is used first.

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 Protection** section, select **Offload to IBM Spectrum Protect** to enable offloading of IBM Spectrum Protect Plus restore points to IBM Spectrum Protect vStorage backup servers.

Offloading essentially creates two backups of your data – one on the vSnap server for short term protection, and one on the IBM Spectrum Protect server for longer term protection. Enter the backup frequency and interval through the associated pulldown menus.

If **Leverage most recent backup** is selected, the offload occurs from the ESX original host or cluster directly, and the latest backup image on the vSnap server is mounted. Note that incremental backups are not supported if selected.

Important: Microsoft Hyper-V is not currently supported for offloading.

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 174
Create a backup job definition that uses the SLA policy.	See the backup topics in Chapter 7, “Protecting hypervisors,” on page 71 and Chapter 8, “Protecting applications,” on page 101.

Related concepts

“Offload to IBM Spectrum Protect ” on page 6

Although the vSnap server is the primary backup location for snapshots, you can also offload snapshots to IBM Spectrum Protect.

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.

Chapter 7. 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 21.

Support for VMware tags

IBM Spectrum Protect Plus supports VMware virtual machine tags. Tags are applied in vSphere and allow users to assign metadata to virtual machines. When applied in vSphere and added to the IBM Spectrum Protect Plus inventory, virtual machine tags can be viewed through the **View > Tags & Categories** filter when you create a job definition. For more information about VMware tagging, see [Tagging Objects](#).

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 to an alternate location, the encrypted virtual machine is restored without encryption, and must be encrypted manually through vCenter Server after the restore completes.

The following vCenter Server privileges are required to enable operations for encrypted virtual machines:

- Cryptographer.Access
- Cryptographer.AddDisk
- Cryptographer.Clone

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 > Backup**.
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. Optional: Use the **IBM Spectrum Protect vStorage Backup Server Settings** to configure an associated vStorage Backup Server for IBM Spectrum Protect offload functionality, select **Link to IBM Spectrum Protect**, and complete the following fields:

Important: The user who registers the vStorage Backup Server must have the **Log on as a service** right enabled to run remote commands. For more information about this right, see [Add the Log on as a service Right to an Account](#).

vStorage Backup Server

Enter the location of the system where the IBM Spectrum Protect for Virtual Environments client GUI and VMCLI are installed.

OS Type

Select the operating system type of the vStorage Backup Server. Available options include Windows and Linux.

Use existing user

Enable to select a previously entered user name and password for the vStorage Backup Server.

vStorage Backup Server Username

Enter your user name for the vStorage Backup Server.

vStorage Backup Server Password

Enter your password for the vStorage Backup Server.

Offloading creates two backups of your data: one on the vSnap server for short-term protection, and one on the IBM Spectrum Protect storage server for long-term protection. Once configured here, offloading is enabled through SLA policies. The offload operation uses data movers from IBM Spectrum Protect for Virtual Environments configured nodes.

7. 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 174 .

Related concepts

[“Offload to IBM Spectrum Protect ” on page 6](#)

Although the vSnap server is the primary backup location for snapshots, you can also offload snapshots to IBM Spectrum Protect.

[“Managing identities” on page 179](#)

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 75](#)

Use a backup job to back up VMware data including virtual machines, datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data” on page 82](#)

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
Cryptographic Operations	<ul style="list-style-type: none"> • Add disks • Direct access • Encrypt • Encrypt new • Manage encryption policies
Datastore	<ul style="list-style-type: none"> • Allocate space • Browse datastore • Low level file operations • Remove file • Rename datastore • Update virtual machine files
Folder	<ul style="list-style-type: none"> • Create folder • Delete folder
Global	<ul style="list-style-type: none"> • Cancel task
Host > Configuration	<ul style="list-style-type: none"> • Storage partition configuration

vCenter Server object	Required privileges
Inventory Service > vSphere Tagging	<ul style="list-style-type: none"> • Assign or Unassign vSphere Tag • Create vSphere Tag • Create vSphere Tag Category • Delete vSphere Tag • Delete 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 VM • Migrate powered on VM • Query vMotion
Virtual Machine > Configuration	<ul style="list-style-type: none"> • Add existing disk • Add new disk • Add or remove device • Advanced • Change CPU count • Disk change tracking • Memory • Modify device settings • Query Fault Tolerance compatibility • Raw device • Reload from path • Remove disk (detach and remove virtual disk) • Rename • Settings
Virtual Machine > Interaction	<ul style="list-style-type: none"> • Backup operation on virtual machine • Power Off • Power On
Virtual Machine > Inventory	<ul style="list-style-type: none"> • Register • Remove • Unregister
Virtual Machine > Provisioning	<ul style="list-style-type: none"> • Allow disk access • Clone virtual machine • Mark as template • Mark as virtual machine

vCenter Server object	Required privileges
Virtual Machine > Snapshot management	<ul style="list-style-type: none"> • Revert to snapshot
vApp	<ul style="list-style-type: none"> • Add VM to vApp • Assign resource pool to vApp • Assign vApp to another 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 > Backup**.
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 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 > Backup**.
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 data including virtual machines, datastores, folders, vApps, and datacenters with snapshots.

Before you begin

Review the following procedures and considerations before you create a backup job definition:

- Register the providers that you want to back up. For more instructions, see [“Adding a vCenter Server instance”](#) on page 71.
- Configure SLA policies. For more instructions, see [“Create backup policies”](#) on page 52.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles must be assigned to the user. Grant users access to hypervisors and backup and restore operations by using the **Accounts** pane. Roles and associated permissions are assigned during user account creation. For more information, see [Chapter 13, “Managing user access,”](#) on page 169 and [“Managing user accounts”](#) on page 177.
- 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.
- A virtual machine is decommissioned when it expires based on the SLA Policies’ retention parameters.
- 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.](#)

Procedure

To define a VMware backup job, complete the following steps:

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

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**, and **Tags & Categories**. 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 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, or you can manually run the job by clicking **Jobs and Operations > Monitoring** and starting the job from the **Job Monitor** pane.

Once 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, you can exclude individual VMDKs from the SLA policy, see [“Excluding VMDKs from the SLA policy for a job”](#) on page 79.

5. To edit options before you create the job definition, click **Select Options**. Set the following job definition options:

Make VM snapshot application/file system consistent

Enable 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 related to quiesced applications.

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.

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 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.

Note: File indexing and file restore are not supported from restore points that were offloaded to IBM Spectrum Protect storage.

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 (UID) 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 offloaded to IBM Spectrum Protect storage.

VM Snapshot retry attempts

Set the number of times that IBM Spectrum Protect Plus attempts to snapshot a virtual machine before the job is canceled.

VADP Proxy

Select a VADP proxy for load sharing and load balancing.

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, then click **Select Options**. Select **Use existing user** and select a previously entered user name and password for the resource. The **Test** button displays to the right of the **Save** button in the **Options** section. Click **Test**.

7. Click **Save**.

8. To configure additional options, click the **Options** field that is associated with the job in the **SLA Policy Status** section. If no additional options are configured for the job, **Not Configured** is shown in the field. Set the additional job options:

Pre-scripts and Post-scripts

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 by using the **Jobs and Monitoring > Script** pane.

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 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 Base 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 backup proxies” on page 80 .
Create a VMware restore job definition.	See “Restoring VMware data” on page 82 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 144](#)

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 **Scripts** pane, and then applied to job definitions.

Related tasks

[“Starting jobs” on page 144](#)

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.

Procedure

To exclude VMDKs from the SLA policy:

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware > Backup**.
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}: 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}
```

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. 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.

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 175](#).

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.

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.

Creating VADP backup proxies

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Before you begin

Note the following considerations before creating VADP proxies:

- Review the IBM Spectrum Protect Plus system requirements in [“VADP proxy requirements” on page 18](#).
- 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 username 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.

What to do next

After you create the VADP proxies, complete the following actions:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 75 . The proxies are indicated in the job log by a log message similar to the following text: Run remote vmdkbackup of MicroService: http://<proxy nodename, IP:proxy_IP_address
Uninstall the proxies when you cease running the VMware backup jobs.	To uninstall a proxy, run uninstall_vmdkbackup from the following directory on the host system: /opt/IBM/SPP/uninstall

Related tasks

[“Setting options for VADP backup proxies” on page 81](#)

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Setting options for VADP backup proxies

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

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 options icon ******* to view available options for the proxy.
3. 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. To enable automatic updates of the VADP proxy, a previously entered user name must be selected.

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 VMDK tasks the proxy can process. If **DEFAULT** is selected, the number of available CPU cores will determine the task limit through the following formula: 1 core = 1 task slot to CPU x number of cores = number of task slots.

What to do next

After you create the VADP proxies, complete the following actions:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 75. The proxies are indicated in the job log by a log message similar to the following text: Run <code>remote vmdkbackup of MicroService: http://<proxy nodename, IP:proxy_IP_address</code>
Uninstall the proxies when you cease running the VMware backup jobs.	To uninstall a proxy, run the following command on the host system from the <code>uninstall</code> subdirectory of the installation directory <code>/opt/IBM/SPP</code> : <code>./uninstall_vmdkbackup</code>

Related tasks

[“Creating VADP backup proxies”](#) on page 80

You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

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

Note the following procedures and considerations before creating a restore job definition:

- Create and run a VMware backup job. For instructions, see [“Backing up VMware data”](#) on page 75.
- Before an IBM Spectrum Protect Plus user can complete backup and restore operations, roles must be assigned to the user. Grant users access to hypervisors and backup and restore operations through the **Accounts** pane. Roles and associated permissions are assigned during user account creation. For more information, see [Chapter 13, “Managing user access,”](#) on page 169 and [“Managing user accounts”](#) on page 177.
- The size of a virtual machine that is restored from a vSnap offload 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 offload. The full size of the data must be transferred even if it is unallocated in the source virtual machine.

- When selecting virtual machines for recovery, restore points offloaded to the IBM Spectrum Protect server cannot be recovered in test mode.
- Restore points that were offloaded to the IBM Spectrum Protect server cannot be used to recover VMDKs.
- Before you select a destination for a restore job definition, ensure that the destination is registered in IBM Spectrum Protect Plus. This includes restore jobs that restore data to original hosts or clusters.

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 after testing and verification are completed. 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 through 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 86.

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 through 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 this 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 recover 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.

You can also 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.

Procedure

To define a VMware restore job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware > Restore**.
2. In the **Restore** pane, review the available restore points of your VMware sources, including virtual machines, virtual machine 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. To select the latest restore point, complete the following actions:
 - a) Click the add to restore list icon  at the resource level.

- b) From the **Restore from site** list, select the site associated with the backup storage server that you want to recover from.

To restore a specific restore point from a specific site, complete the following actions:

- a) Expand a resource in the **Restore** pane.
- b) Click the add to restore list icon  that is associated with the restore point.

Adding the latest restore points and specific restore points to the **Restore List** is not supported. Click the remove icon  to remove restore points from the **Restore List**.

4. To run the job now with the default options, click **Restore**. To schedule the job to run with the default options, click **Manage Jobs** and define a trigger for the job definition.
5. To edit options, click **Options**. You can set the following job definition options:

Destination

Set the VMware destination:

Original ESX Host or Cluster

Select to restore to the original host or cluster.

Alternate ESX Host or Cluster

Select to restore to a local destination different from the original host or cluster, and then select the alternate location from 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 vCenter section, select an alternative location. Selections can be filtered by either hosts or clusters.

ESX host if vCenter is down

Select to bypass the vCenter and restore directly to the ESX host. In other restore scenarios, actions are completed through vCenter. If vCenter is unavailable, this option restores the vCenter virtual machine or virtual machines that the vCenter is dependent on.

Restore Type

Set the VMware restore job to run in test, production, or clone mode by default. After the job is created, it can be run in production or clone mode through the Job Sessions or Active Clones sections of the **Restore** pane. Test mode is not available for long-distance restore operations.

Network Settings

Set the network settings for a restore to an original ESX host or cluster:

Allow system to define IP configuration

Select 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 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.

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 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 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 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.

Script Settings

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 through the **Jobs and Monitoring > Script** pane.

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.

Advanced Options

Set the advanced job definition options:

Power on after recovery

Toggle the power state of a virtual machine after a recovery is performed. Virtual machines are powered on in the order they are recovered, as set in the Source step. Note that restored virtual machine templates cannot be powered on after recovery.

Overwrite virtual machine

Enable 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 to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up 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.

Fix VMX file for missing disk

If individual disks are excluded from a backup, the associated virtual machine will fail to start. Enable this option to remove the entries for excluded disks from the VMX configuration file and ensure the restored virtual machine starts as part of an Instant VM restore job.

Restore VM tags

Enable this option to restore tags applied to virtual machines through vSphere.

Click **Save** to save the policy options.

6. To run the job now, click **Restore**. To schedule the job click **Manage Jobs** and define a trigger for the job definition.
7. 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.

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 71](#)

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 82](#).
- When selecting virtual machines for recovery, restore points offloaded to the IBM Spectrum Protect server cannot be recovered in Test Mode.
- Restore points that were offloaded to the IBM Spectrum Protect server cannot be used to recover VMDKs.

Procedure

To create a fenced network, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware > Restore**.
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 71](#)

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 or other management VMs are not accessible

IBM Spectrum Protect Plus provides an option to automatically restore data by using ESXi hosts if the vCenter is not accessible. This option restores the vCenter virtual machine (VM) or VMs that the vCenter is dependent on.

About this task

This procedure can be used if any of the following management services are partially or fully lost in your environment:

- vCenter
- Platform Services Controller (PSC)
- Software-Defined Data Center (SDDC)
- Active Directory (AD)
- Domain Name System (DNS) servers

To recover data without a vCenter, the ESXi host must have a standard switch or a preexisting distributed switch with ephemeral binding. If these requirements are not met, you must create a new standard switch on the ESXi host. If there are no available uplinks for the standard switch, the standard switch must be removed from the distributed switch.

The procedure describes additional manual steps are required to complete a restore operation when the operation is running in a vCenter Server (VCS) environment.

Recovering a management VM in a VCS environment can result in the loss of access to the VM. The loss of access is due to a misconfiguration of the virtual switch. Complete the following steps on the affected VM to recover from this state and to complete a recovery operation.

Procedure

1. Connect to the destination ESXi user interface host and create a new standard virtual switch. At this point, there are no port groups or uplinks available for the switch.
2. Use the SSH protocol to connect to the ESXi server. Identify and select the physical NIC and the port group of the existing distributed virtual switch that is named SDDC-Dswitch-Private. The following example references a virtualized Network Interface Card (vNIC) named `vmnic0`, which is part of port ID 64. You can list the Distributed Virtual Switch (DVS) information by issuing the following command:

```
#esxcli network vswitch dvs vmware list
```

3. Based on the previous information, remove the NIC and port ID (port group) from the SDDC-Dswitch-Private DVS by using the following command. Use the port ID from step 2.

```
#esxcfg-vswitch -Q physical_vnic -V port_group SDDC-Dswitch-Private
```

4. Add the NIC and port group into the standard switch that you created in step 1 by issuing the following command on one line:

```
#esxcli network vswitch standard uplink add --uplink-name=physical_vnic --vswitch-name=standard_vswitch
```

5. In the ESXi interface, add a port group and select the standard virtual switch.
The virtual switch should have one uplink and one port group.
6. Run a restore operation in IBM Spectrum Protect Plus with the **ESX host if vCenter is down** option enabled.
7. Click **Options** when defining the restore operation in IBM Spectrum Protect Plus and choose the new networking switch that you created in step 1 under **Networking**.
8. Using the destination ESXI user interface, power on the recovered VM.
9. After the VMs are reachable, log in to the vCenter user interface and start the migration of the management VMs from the temporary port group that you created in step 5 to the original distributed port group, SDDC-DPortGroup-Mgmt.

Start a migration from the **Networking** tab by selecting a datacenter and then clicking **Migrate VMs to Another Network** on the **Actions** menu. Select the source network (the temporary switch created in step 5) and the destination network (the management switch).

10. After all VMs are migrated into the original port group, reincorporate the physical NIC and the port group into the original distributed virtual switch by taking the following actions:
 - a. Remove the network cards (known as vmnics) from a standard vSwitch that was reassigned previously 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=veried_recovery
```

- b. Add network cards to a vNetwork Distributed Switch (vDS) by issuing the following command:

```
#esxcfg-vswitch -P vmnic -V unused_dvPort_ID dvSwitch # add a vDS uplink
```

For example:

```
#esxcfg-vswitch -P vmnic0 -V 64 SDDC-Dswitch-Private
```

11. Delete the temporary port group and the standard vSwitch from the ESXi host user interface.
12. 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 unregisters a VM from the vCenter catalog but leaves VMDK files on the datastore, which consume storage space, on the datastore. 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 21.

Adding a Hyper-V server

When a Hyper-V server is added to IBM Spectrum Protect Plus, an inventory of the server is captured, enabling you to complete backup and restore jobs, as well as run reports.

Before you begin

Note the following considerations and procedures before adding a Hyper-V server to IBM Spectrum Protect Plus:

- Hyper-V servers can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Protect Plus. If the Hyper-V server is part of a cluster, all nodes in the cluster must be resolvable through DNS. If DNS is not available, the server must be added to the `/etc/hosts` file on the IBM Spectrum Protect Plus appliance. If more than one Hyper-V server is set up in a cluster environment, all of the servers must be added to `/etc/hosts`. When registering the cluster in IBM Spectrum Protect Plus, register the Failover Cluster Manager.
- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI initiator Service running in their Services list. Set the service to Automatic so that it is available when the machine boots.
- Add the user to the local administrator group on the Hyper-V server.

Procedure

To add a Hyper-V server, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > Hyper-V > Backup**.
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 SLL 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 152.

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 91.

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 174 .

Related tasks

[“Backing up Hyper-V data” on page 92](#)

Use a backup job to backup Hyper-V data with snapshots.

[“Restoring Hyper-V data” on page 95](#)

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 > Backup**.
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 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 > Backup**.
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 backup Hyper-V data with snapshots.

Before you begin

Note the following procedures and considerations before creating a backup job definition:

- Register the providers that you want to back up. For more information see [“Adding a Hyper-V server” on page 90](#)
- Configure SLA policies. For instructions, see [“Create backup policies” on page 52](#).
- 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 2016](#).

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 complete backup and restore operations, roles must be assigned to the user. Grant users access to hypervisors and backup and restore operations through the **Accounts** pane. Roles and associated permissions are assigned during user account creation. For more information, see [Chapter 13, “Managing user access,” on page 169](#) and [“Managing user accounts” on page 177](#).
- 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.
- A virtual machine is decommissioned when it expires based on the SLA Policies’ retention parameters.

Procedure

To define a Hyper-V backup job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Hypervisors > Hyper-V > Backup**.
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 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, or you can manually run the job by clicking **Jobs and Operations > Monitoring** and starting the job from the **Job Monitor** pane.
5. To edit options before creating the job definition, click **Select Options**. Set the following job definition options:

Make VM snapshot application/file system consistent

Enable to turn on application or file-system consistency for the virtual machine snapshot.

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.

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 through the **File Restore** pane in IBM Spectrum Protect Plus. 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. Ensure that the virtual machine can be accessed from the IBM Spectrum Protect Plus appliance either through DNS or host-name.

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 (UID) 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 offloaded to an IBM Spectrum Protect server.

VM Snapshot retry attempts

Set the number of times IBM Spectrum Protect Plus should attempt to snapshot a virtual machine before canceling the job.

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. The **Test** button displays to the right of the **Save** button in the **Options** section. Click **Test**.

7. Click **Save**.

8. To configure additional options, click the **Options** field that is associated with the job in the **SLA Policy Status** section. If no additional options are configured for the job, **Not Configured** is shown in the field. Set the additional job options:

Pre-scripts and Post-scripts

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 through the **Jobs and Monitoring > Script** pane.

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 Base 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 95 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 144](#)

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 **Scripts** pane, and then applied to job definitions.

Related tasks

[“Starting jobs” on page 144](#)

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

Note the following procedures and considerations before creating a restore job definition:

- Create and run a Hyper-V Backup job. For instructions, see [“Backing up Hyper-V data” on page 92](#).
- When selecting a destination for a restore job definition, note that the destination must be registered in IBM Spectrum Protect Plus. This includes Restore jobs that restore data to original hosts or clusters.
- 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 2016](#).

For Linux environments, see [Supported Linux and FreeBSD virtual machines for Hyper-V on Windows](#).

- Before an IBM Spectrum Protect Plus user can complete backup and restore operations, roles must be assigned to the user. Grant users access to hypervisors and backup and restore operations through the **Accounts** pane. Roles and associated permissions are assigned during user account creation. For more information, see [Chapter 13, “Managing user access,” on page 169](#) and [“Managing user accounts” on page 177](#).

About this task

If a Virtual Hard Disk (VHDX) is selected for restore, 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, 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 after testing and verification are completed. 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 created through 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 requiring permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines created through 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 recover 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.

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 > Restore**.
2. In the **Restore** pane, review the available restore points of your Hyper-V 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. To select the latest restore point, complete the following actions:
 - a) Click the add to restore list icon  at the resource level.
 - b) From the **Restore from site** list, select the site associated with the backup storage server that you want to recover from.

To restore a specific restore point from a specific site, complete the following actions:

- a) Expand a resource in the **Restore** pane.
- b) Click the add to restore list icon  that is associated with the restore point.

Adding the latest restore points and specific restore points to the **Restore List** is not supported. Click the remove icon  to remove restore points from the **Restore List**.

4. To run the job now with the default options, click **Restore**. To schedule the job to run with the default options, click **Manage Jobs** and define a trigger for the job definition.
5. To edit options before creating the job definition, click **Options**. Set the job definition options.

Destination

Set one of the following Hyper-V destinations:

Original Hyper-V Host or Cluster

Select to restore to the original host or cluster.

Alternate Hyper-V Host or Cluster

Select to restore to a local destination different from the original host or cluster, then select the alternative location from available resources.

Restore Type

Set the Hyper-V restore job to run in test, production, or clone mode by default. After the job is created, it can be run in test, production, or clone mode by using the **Job Sessions** pane.

Network Settings

Set the network settings for a restore to an alternate Hyper-V 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.

Destination Datastore

Set the destination datastore for a restore to an alternate Hyper-V 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.

Script Settings

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 through the **Jobs and Monitoring > Script** pane.

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.

Advanced Options

Set the following advanced job definition options:

Power on after recovery

Toggle the power state of a virtual machine after a recovery is performed. Virtual machines are powered on in the order they are recovered, as set in the Source step. Note that restored VM templates cannot be powered on after recovery.

Overwrite virtual machine

Enable 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 to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up 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.

Click **Save** to save the policy options.

6. To run the job now, click **Restore**. To schedule the job click **Manage Jobs** and define a trigger for the job definition.
7. 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 92](#)

Use a backup job to backup Hyper-V data with snapshots.

[“Adding a Hyper-V server” on page 90](#)

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 21](#).
- Run a backup job with catalog file metadata enabled. Follow these guidelines:
 - Ensure that credentials are established for the associated virtual machine as well as the alternate virtual machine destination through the Guest OS Username and Guest OS Password option within the backup job definition.
 - Ensure that the virtual machine can be accessed from the IBM Spectrum Protect Plus appliance either through DNS or hostname. In a Windows environment, the default security policy uses the Windows NTLM protocol, and the user identity follows the default *domain\name* format if the Hyper-V virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.
 - For a file restore to complete successfully, ensure that the user ID that is on the target machine has the necessary ownership permissions for the file that is being restored. If a file was created by a user that differs from the user ID that is restoring the file based on Windows security credentials, the file restore job fails.

About this task

Restrictions:

- Encrypted Windows file systems are not supported for file cataloging or file restore.
- File indexing and file restore are not supported from restore points that were offloaded to IBM Spectrum Protect 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.

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

Procedure

To restore a file, complete the following steps.

1. In the navigation pane, click **Jobs and Operations > 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 183.
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 75](#)

Use a backup job to back up VMware data including virtual machines, datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data” on page 82](#)

VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Chapter 8. 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.

Db2

After you successfully add your IBM Db2 instances to **Manage Protection > Applications > Db2**, you can start to protect your Db2 data with IBM Spectrum Protect Plus. Once this is complete, create jobs to back up your Db2 data and define retention policies to maintain the data.

Prerequisites for Db2

All prerequisites for IBM Spectrum Protect Plus Db2 application server must be in place before you start protecting Db2 data with IBM Spectrum Protect Plus.

The 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 operations, and in the target volumes for the copying of files during the restore operation. For more information about space requirements, see [Space requirements for Db2 protection](#). When you are restoring data to an alternate location, you must 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, but the paths are on separate volumes. This setup is required to allow copying of data from the mounted vSnap to the target host. Ensure that you have dedicated local database directories for each database in your volume setup.

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.
- If you are running IBM Spectrum Protect Plus in an AIX or Linux environment, check the version of sudo installed to ensure it is at the recommended level as detailed in the [Prerequisites](#) document for your version.
- Ensure the Linux `util-linux-ng` or `util-linux` package is up-to-date.
- Ensure that the database table spaces, online logs, and the local database directory are located on dedicated logical volumes as shown in the following figure. One layout shows all volumes for data and logs on one volume group, the other shows 2 volume groups.

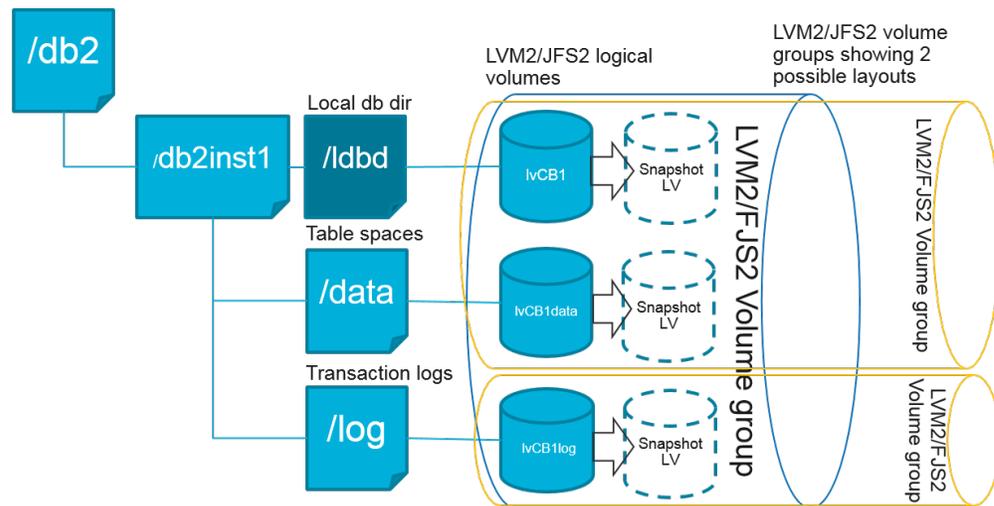


Figure 10: Logical volume layout examples

- Ensure that you do not have 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 at the vSnap repository. Extra free disk space is required on the volume groups for creating temporary LVM snapshots of the logical volumes that the Db2 database and logs files are stored on. To create LVM snapshots of a protected Db2 database, ensure that there is enough free disk space on the volume groups with Db2 data.

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 temporarily create a consistent point-in-time copy of the Db2 database.

When the IBM Spectrum Protect Plus Db2 agent creates the LVM snapshot, it is mounted and a file copy operation to the vSnap repository starts. 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.

Note:

- On Linux, LVM snapshots are created with Linux `lv` commands.
- On AIX, LVM snapshots are created with the JFS2 snapshot command as external JFS2 snapshots.

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 there are no internal snapshots on the JFS2 volumes as these 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 there is enough free disk space in the volume group, 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. On AIX, file systems the logical volumes are managed by JFS2.

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 Linux, the LVM2 volume manager stores the snapshot of a logical volume within the same volume group. On AIX, 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 with data as it changes on the source volume while the snapshot exists.

Setting sudo privileges for Db2

To use IBM Spectrum Protect Plus to protect your data, you must install the required sudo version. 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

Review the [Prerequisites](#) for the correct version of sudo to be installed for your version of IBM Spectrum Protect Plus.

About this task

Set up a dedicated IBM Spectrum Protect Plus agent user with the required super user privileges with sudo. This configuration enables sudoers to run commands without a password.

Procedure

1. Create an application server user.
For example, `useradd -m <agent>`
2. Set a password for the new user.
For example, `passwd <agent>`
3. Set the `!requiretty` setting for your system.
4. At the end of the sudoers configuration file `/etc/sudoers`, add these lines.

```
Defaults:agent !requiretty
agent ALL=(ALL) NOPASSWD:ALL
```

If your sudoers file is configured to import configurations from another directory `/etc/sudoers.d`, you can add the lines in a new file in that directory.

Adding a Db2 application server

To start protecting your Db2 data, you must add the machine that hosts your Db2 instances. You can repeat the procedure to add all the machines that host Db2 data that you want to protect with 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 > Backup**.
2. In the **Backup** window, click **Manage Application Servers**, and click **Add Application Server** to add the host machine.

Add Application Server

3. In the **Application Properties** section, enter the host address.
4. Either select an existing user, or enter a user ID and password.
The user must have sudo privileges set up.

The screenshot shows the IBM Spectrum Protect Plus web interface. On the left is a dark navigation pane with a home icon and menu items: Dashboard, Jobs and Operations, Manage Protection, Policy Overview, Hypervisors, Applications, Db2, Backup, Restore, Oracle, SQL, Backup Spectrum Protect Plus, and Restore Spectrum Protect Plus. The 'Backup' menu item is highlighted. The main content area is titled 'Backup' and contains a 'Manage Application Servers' button. Below this is a 'Manage Application Servers' header and an 'Application Properties' form. The form has four fields: 'Host Address' with the value 'ac.ie.ibm.com/dblsntf', 'Use existing user' with an unchecked checkbox, 'UserId' with the value 'spp_ac.ie.ibm.com', and 'Password' with masked characters '*****'. At the bottom of the form are 'Cancel' and 'Save' buttons.

5. Save the form, and repeat the steps to add other Db2 application servers to IBM Spectrum Protect Plus.

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, check the job log. To do this, select **Jobs and Operations** in the navigation pane. Click **Monitoring > Job Monitor** and look for the latest Default Application Server Inventory log entry. Databases must be detected to ensure that they can be protected. For instructions about running an on demand inventory, see [Detecting Db2 resources](#).

Detecting Db2 resources

After you add your IBM Db2 application servers to IBM Spectrum Protect Plus, an inventory to detect all Db2 instances and databases is run automatically. You can run an inventory on any application server manually to detect, list and store all Db2 databases for the selected host making them available for protection with IBM Spectrum Protect Plus.

Before you begin

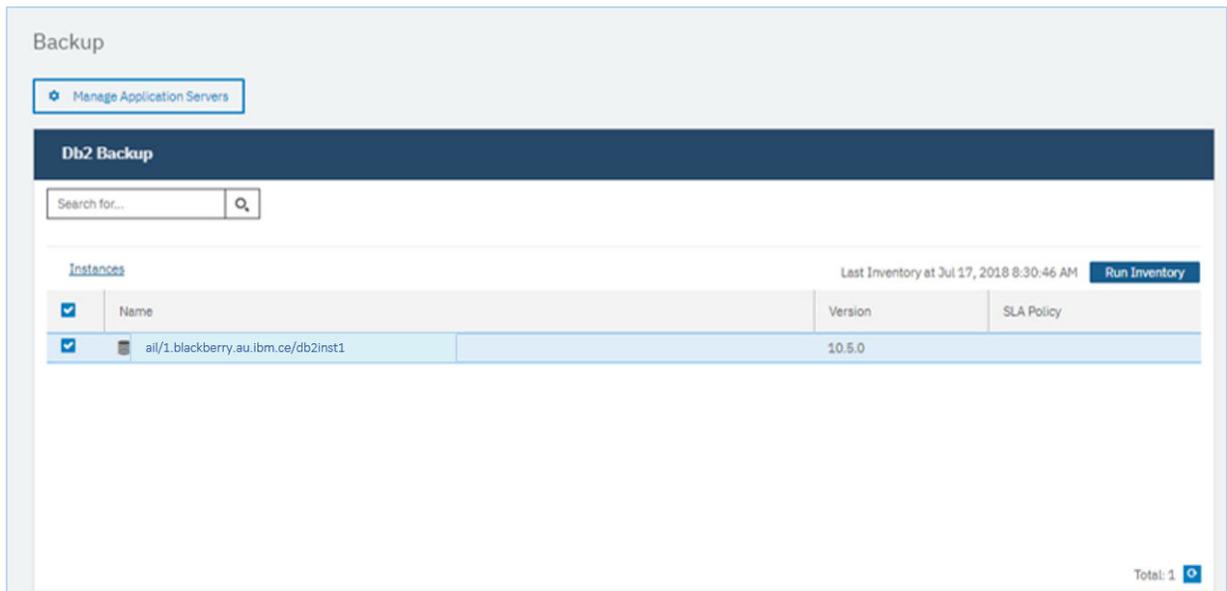
Ensure that you added your Db2 application servers to IBM Spectrum Protect Plus. For instructions, see [Adding a Db2 application server](#).

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > Db2 > Backup**.

Tip: To add more Db2 instances to the **Instances** pane, follow the instructions in [Adding a Db2 application server](#).

2. Click **Run Inventory**.



When the inventory is running, the button changes to show **Inventory In Progress**. You can run an inventory on any available application servers, but you can run only one inventory process at a time.

To monitor the inventory job, select **Jobs and Operations** in the navigation, **Monitoring > Job Monitor**. The job is listed as **Default Application Server Inventory**.

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 an Service Level Agreement (SLA) policy to the instance. For instructions about setting an SLA policy, see [Defining an SLA policy](#).

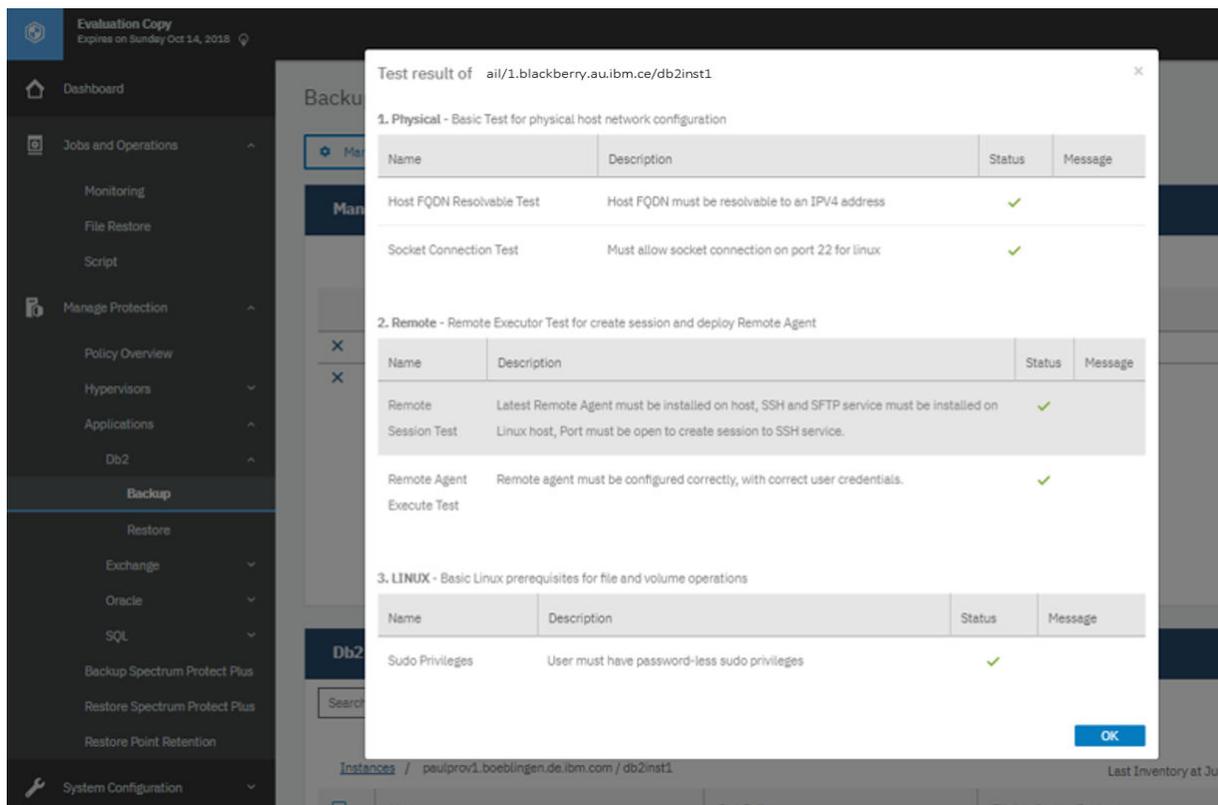
Testing the Db2 connection

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 > Backup**.
2. In the **Backup** 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.



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.

4. Click **OK** to close the test, and choose to rerun the test after you fix any failed tests.

Backing up Db2 data

Define regular Db2 backup jobs with options to run and create backup copies to protect your data. You can enable continuous backing up of archive logs so that you can restore a point-in-time copy with rollforward options if required.

Before you begin

During the initial backup, IBM Spectrum Protect Plus creates a new vSnap volume and NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus Db2 agent mounts the share on the Db2 server where the backup is to be completed.

Review the following procedures and considerations before you create a backup job definition:

- Add the application servers that you want to back up. For the procedure, see [Adding a Db2 application server](#).
- Configure a Service Level Agreement (SLA) Policy. For the procedure, see [Defining a Service Level Agreement backup job](#).
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see [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.

- 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.

Procedure

1. From the navigation , expand **Manage Protection > Applications > Db2 > Backup**.
2. Click the database or instance you want to back up by choosing one of the following choices.
 - Select an entire instance to back up in the **Instances** pane by clicking the check-box beside the instance name. Any databases added later to this instance are automatically adopted to the SLA policy chosen.
 - Select a specific database in an instance by clicking the instance name, and choosing a database from the list of databases in that instance.

Each item in the **Instances** pane, is listed by instance or database name, the applied SLA policy, and the eligibility for log backup.

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. Save 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](#).

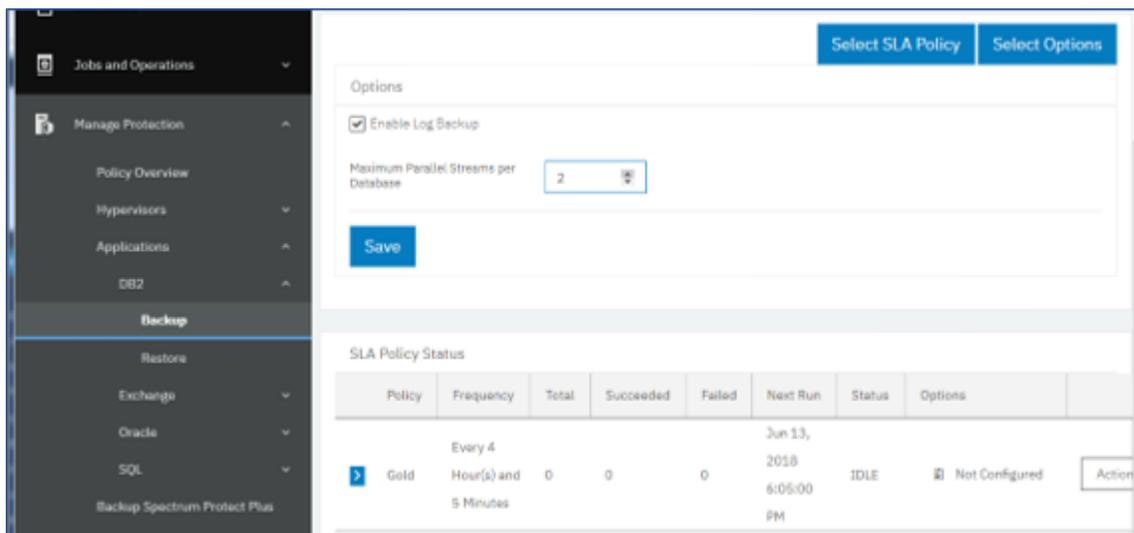


Figure 11: Backup pane that shows Enable Log Backup

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, if you choose to you can run an on-demand backup any time by clicking **Actions** beside the policy name and selecting **Start**. The status in the log changes to show that the backup is Running.

Defining a Service Level Agreement backup job

Once your Db2 databases are listed for each of your Db2 instances, select and apply an Service Level Agreement (SLA) to start protecting your data.

Procedure

1. From the navigation, expand **Manage Protection > Applications > Db2 > Backup**.
2. Select the Db2 instance to back up all the data in that instance, or select individual databases in the Db2 instance that you want to back up.

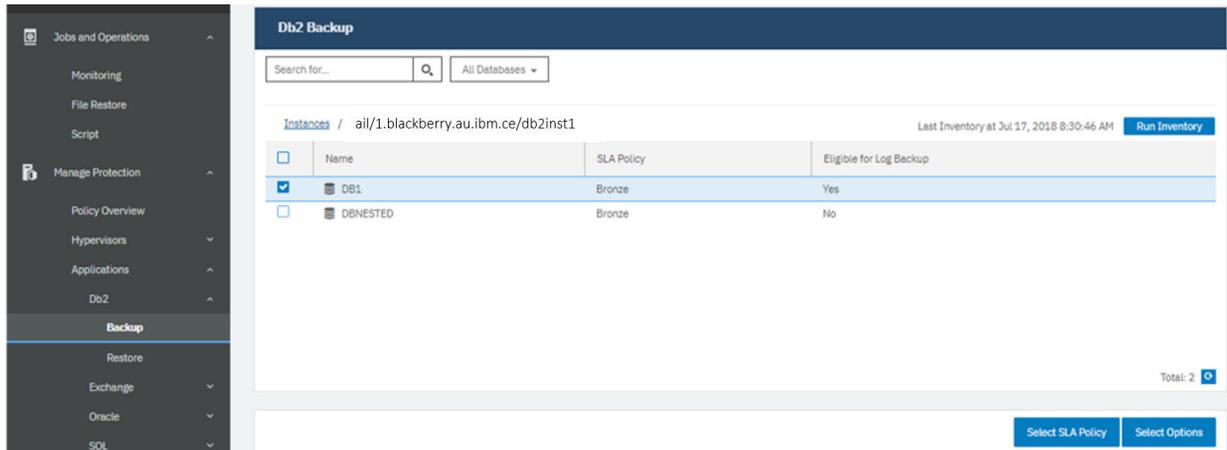


Figure 12: Db2 Backup pane showing instance databases

3. Click **Select 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. 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.

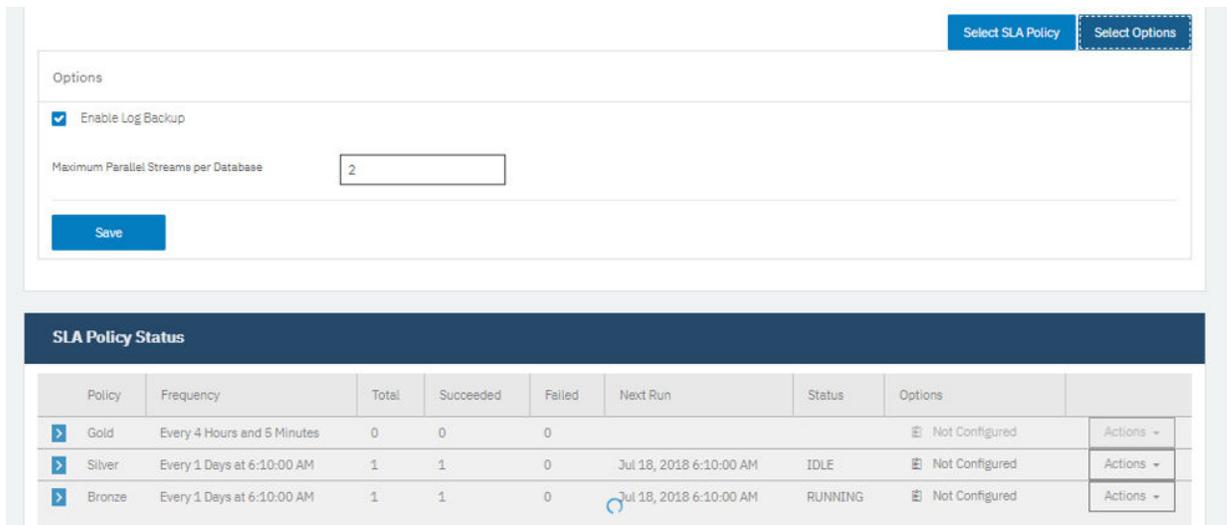


Figure 13: Backup Options and SLA Policies

5. If you want 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.

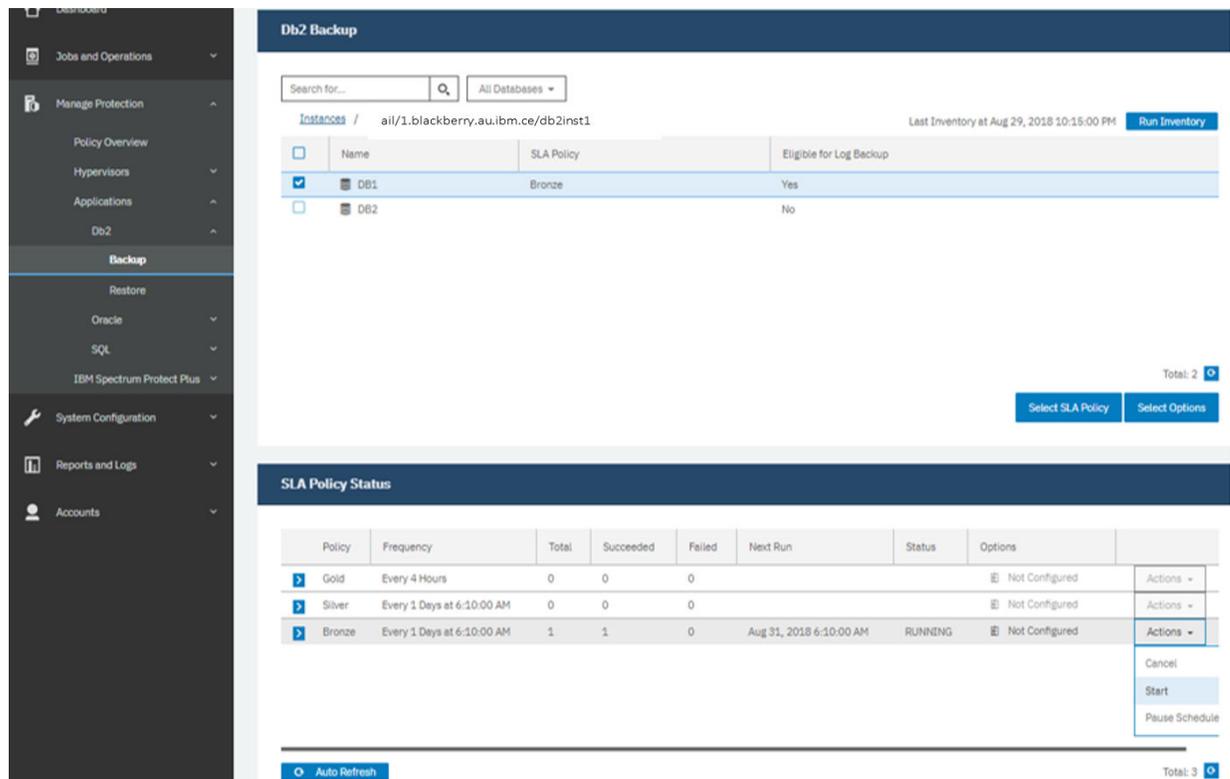


Figure 14: SLA policies

To pause the schedule of an SLA, click **Actions** and choose **Pause Schedule**.

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, and updates either LOGARCHMETH1 or LOGARCHMETH2 parameters to point to that volume for log archiving. The volume is kept mounted on the Db2 application server until the **Enable Log Backup** option is cleared and a new backup job is run.

When either LOGARCHMETH1 or LOGARCHMETH2 parameters are set with a value other than OFF, you can use archived logs for rollforward recovery. You can disable log backup jobs at any time by clearing the **Enable Log Backups** option in **Applications > Db2 > Backup > 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 only enable log backup jobs 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 108](#).

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 data, define a job that restores data from the latest database backup or select an earlier backup copy. Choose to restore data to the original instance or to an alternative instance on a different machine. Define the restore job by specifying recovery options, and save the job.

Before you begin

In the **Restore** pane, use the filters to show the instances from the primary or secondary sites. The default shows the restore points from all sites. If you cannot find a specific database, search by entering a database name in the search field.

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 106](#).
- 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 13, “Managing user access,” on page 169](#).

Before you start a restore operation to an alternate instance, ensure that the same 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 the Db2](#).

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > Db2 > Restore**.
2. In the **Db2 Restore** pane, click a Db2 instance to show the databases in that instance.
3. Expand the database that you want to restore to show the available restore points for that database.

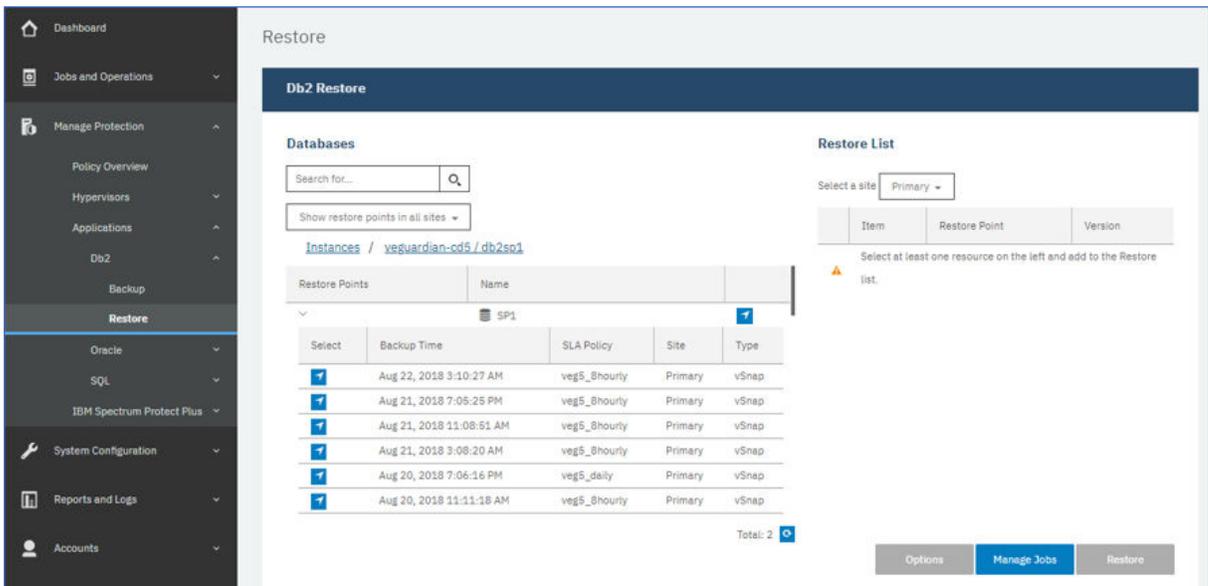


Figure 15: Restore panel for Db2

Restore points are listed with the backup date and time, SLA policy, site information, and type.

4. Choose to restore the latest backup or an earlier backup from the **Restore Points** list:

- To restore the latest backup, click the add icon beside the database name on the right of the **Restore Points** table, .

To restore with recovery to a specific point-in-time, you must select the overall database.

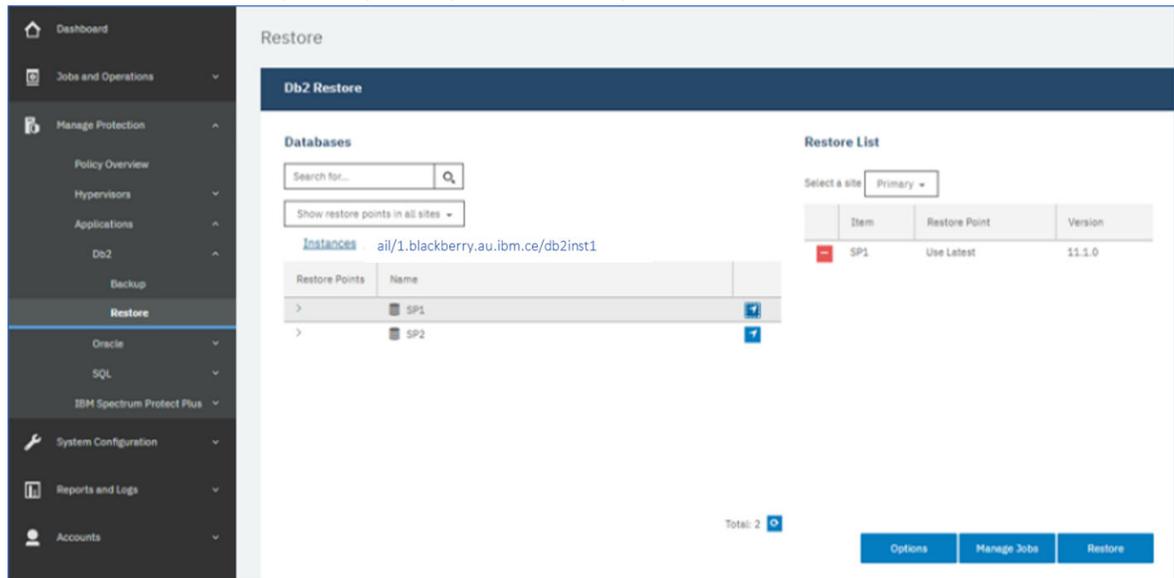


Figure 16: Selecting to add the overall database to the restore list

- To choose a restore point from a different time, find the backup that you require and add it to the restore list by clicking the add icon .

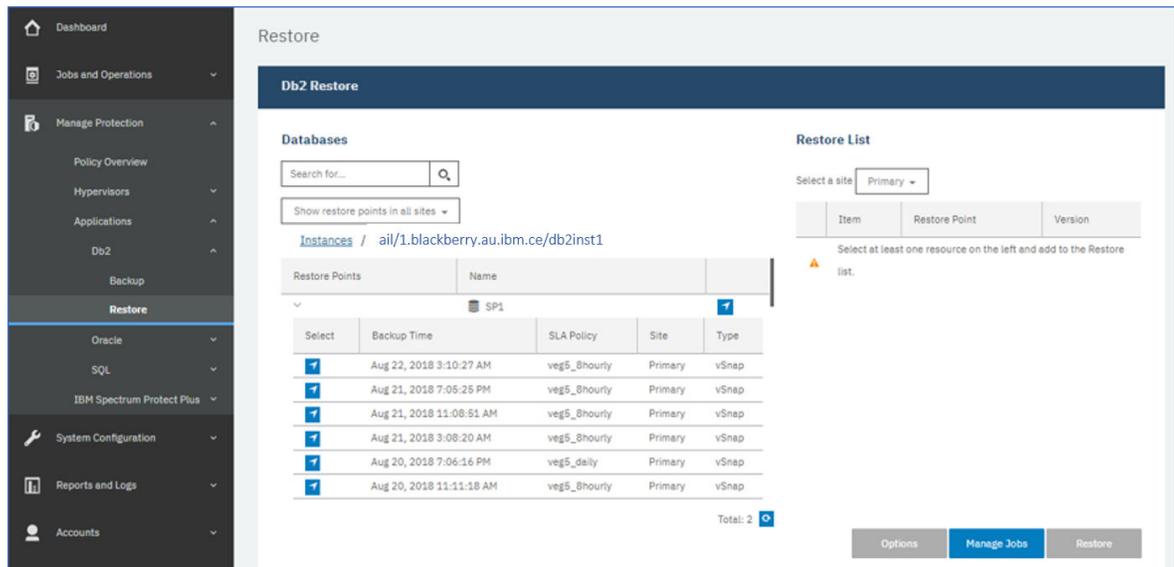


Figure 17: Selecting to add a specific backup the restore list

To remove the restore point from the list, click the delete icon .



Attention: Ensure that you review the selected options before clicking **Restore** because data will be overwritten when the **Overwrite existing data** option is selected.

5. Click **Options** to define specific options for the restore job as follows:

- **Restore Type:** Choose one of the following for your restore operation.
 - **Test:** In this mode, the agent creates a new database by using the data files directly from the vSnap repository. This option is available only when restoring to an alternate instance.
 - **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. This restore type is the only option available when you restore data to the original instance.
 - **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. This option is available only when restoring to an alternate instance.
- **Destination:** Click **Restore to original instance** to restore data to the original server, or **Restore to alternate instance** to restore data to a different location that you can select from the locations listed.

For the procedure to restore data to the original instance, see [Restoring to the original instance](#). For the procedure to restore data to an alternative instance, see [Restoring to an alternate instance](#).

6. Define a recovery option for the restore operation in the **Recover Options** section.

- **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 you have enabled log backups in your Db2 backup job definition. IBM Spectrum Protect Plus uses the newest 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, it is only available when you add the overall database. When you add the overall database, the newest backup is automatically selected for the end of logs recovery.

- **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, May 8, 2018 12:18:00 AM. IBM Spectrum Protect Plus finds the restore points directly before and after the point-in-time chosen. During the recovery process, the older data backup volume and the newer log backup volume are mounted. A temporary restore point is created if the point in time is after the last backup. This recovery option is not available if you selected a specific restore point from the list. Selecting the overall database automatically selects the newest backup for the end of logs recovery process.

7. Select application options in the **Application Options** section as follows:

- **Overwrite existing databases.** This option is selected by default during a Test restore operation to replace existing databases with the same name encountered by the restore process during recovery. If this option is not selected, the restore job fails when databases with the same name are found during the restore process. If you select this option, ensure that the Db2 log directory and the Db2 mirror log directory are clear of data.



Attention: Ensure that no other databases share the same local database directory as the original database or that data is overwritten when this choice is selected.

- **Maximum Parallel Streams per Database.** If required you can choose to run the restore operation of data in parallel streams. This option is useful when 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.

8. Define advanced options for the restore job in the **Advanced Options** section as follows:

- **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.

9. Define **Scripts Settings** as follows:

- Select **Pre-script** to select an uploaded script and an application or script server where the pre-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **Jobs and Monitoring > Script** pane to configure scripts and script servers.
- Select **Post-script** to select an uploaded script and an application or script server where the post-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **Jobs and Monitoring > Script** pane to configure scripts and script servers.
- Select **Continue job/task on script error** to continue running the job when the script associated with the job fails. When this option is enabled, in the event that a script completes with a nonzero return code, the backup or restore job continues to run and the pre-script task status returns COMPLETED. If a post-script completes with a nonzero return code, the Post-script task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task status returns FAILED.

10. Review the job definition and click **Save**.

11. To run the job immediately, click **Restore**. To specify a schedule for a repeated restore operation, click **Manage Jobs** to define a trigger for the job.

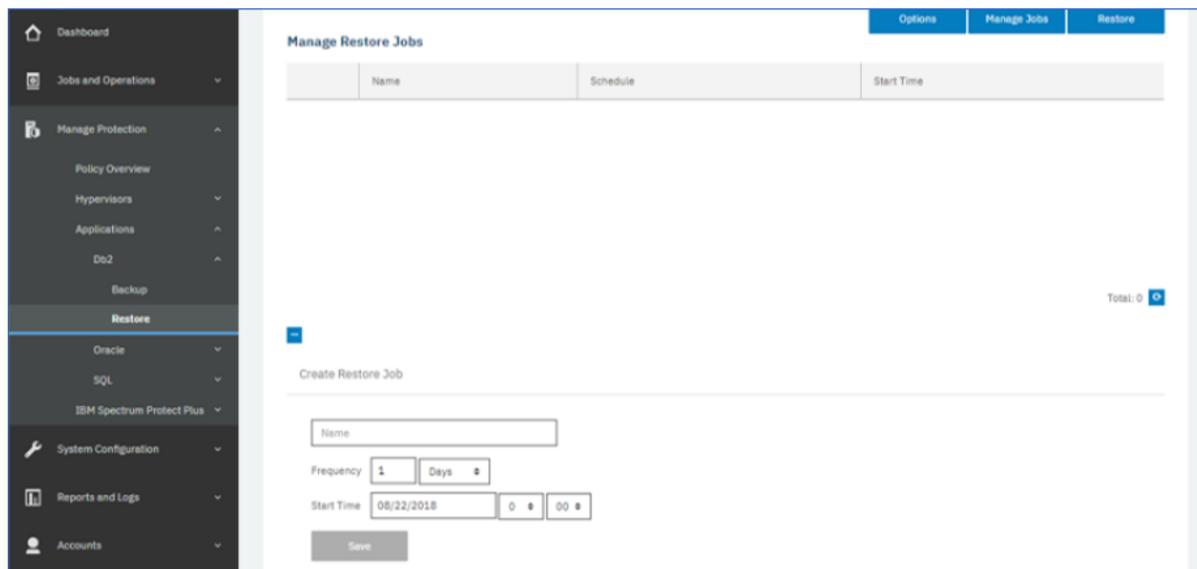


Figure 18: Restore job definition in Manage Restore Jobs

Results

A few moments after you select **Restore**, 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 selecting the download icon .

To restore data to the original instance, follow the instructions in [Restoring to the original instance](#). To restore data to an alternate instance, follow the instructions in [Restoring to an alternate instance](#).

Restoring Db2 data to the original instance

Restore a database backup to its original instance on the original host. Choose between restoring 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

In the **Restore** pane, use the filters to show the instances from the primary or secondary sites. The default shows the restore points from all sites. If you cannot find a specific database, search by entering a database name in the search field.

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 106.
- 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 13, “Managing user access,”](#) on page 169.

Before you start a restore operation to an alternate instance, ensure that the same 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 the Db2](#).

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > Db2 > Restore**.
2. In the **Db2 Restore** pane, click a Db2 instance to show the databases in that instance.
3. Expand the database that you want to restore to show the available restore points for that database.

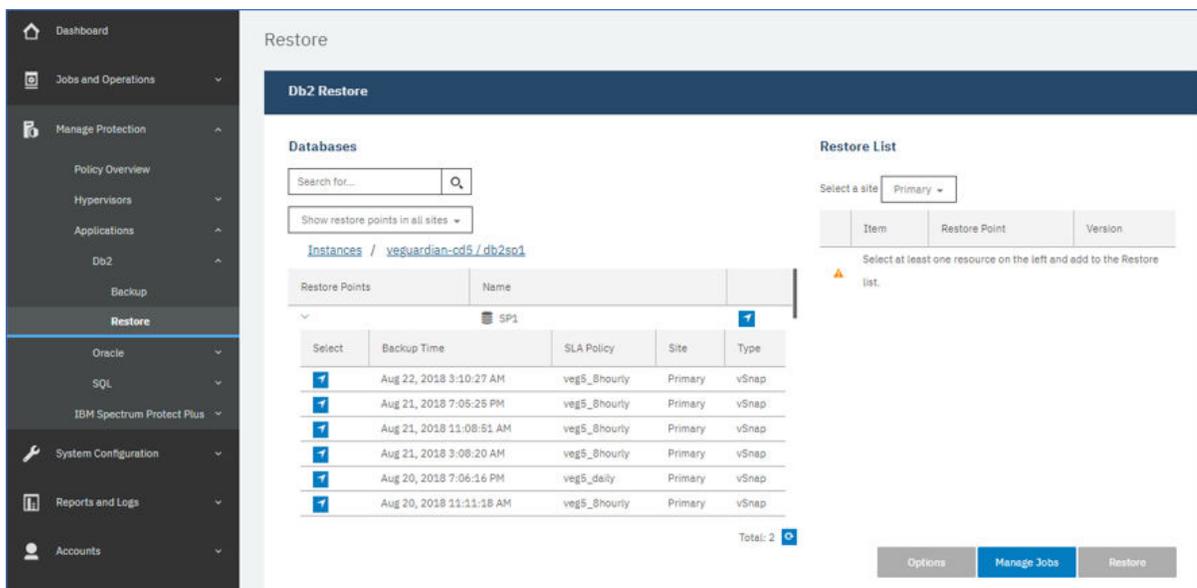


Figure 19: Restore panel for Db2

Restore points are listed with the backup date and time, SLA policy, site information, and type.

4. Choose to restore the latest backup or an earlier backup from the **Restore Points** list:
 - To restore the latest backup, click the add icon beside the database name on the right of the **Restore Points** table, .

To restore with recovery to a specific point-in-time, you must select the overall database.

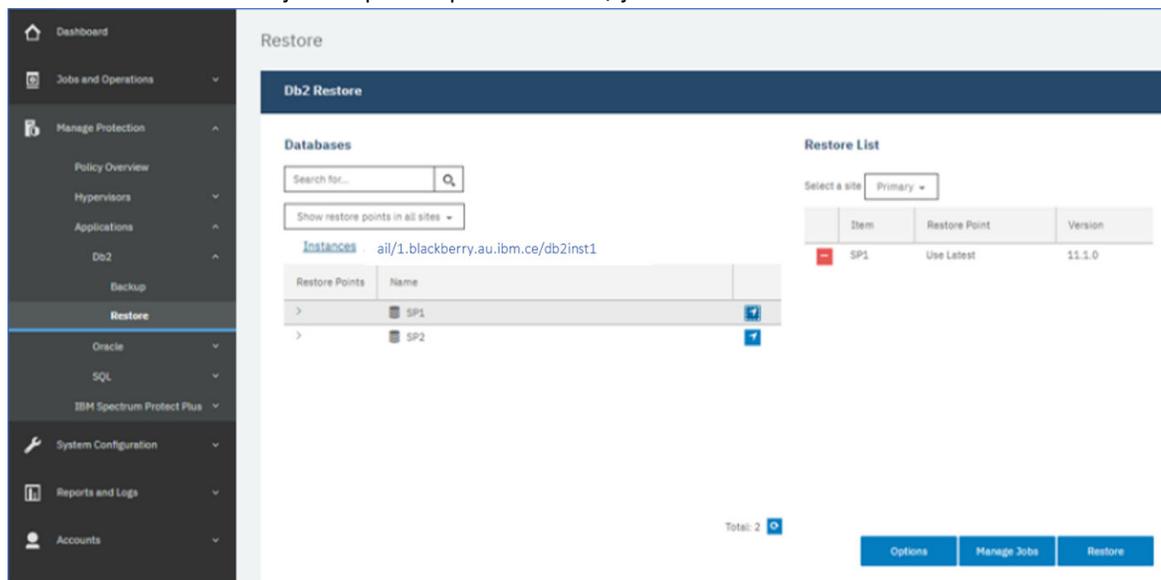


Figure 20: Selecting to add the overall database to the restore list

- To choose a restore point from a different time, find the backup that you require and add it to the restore list by clicking the add icon .

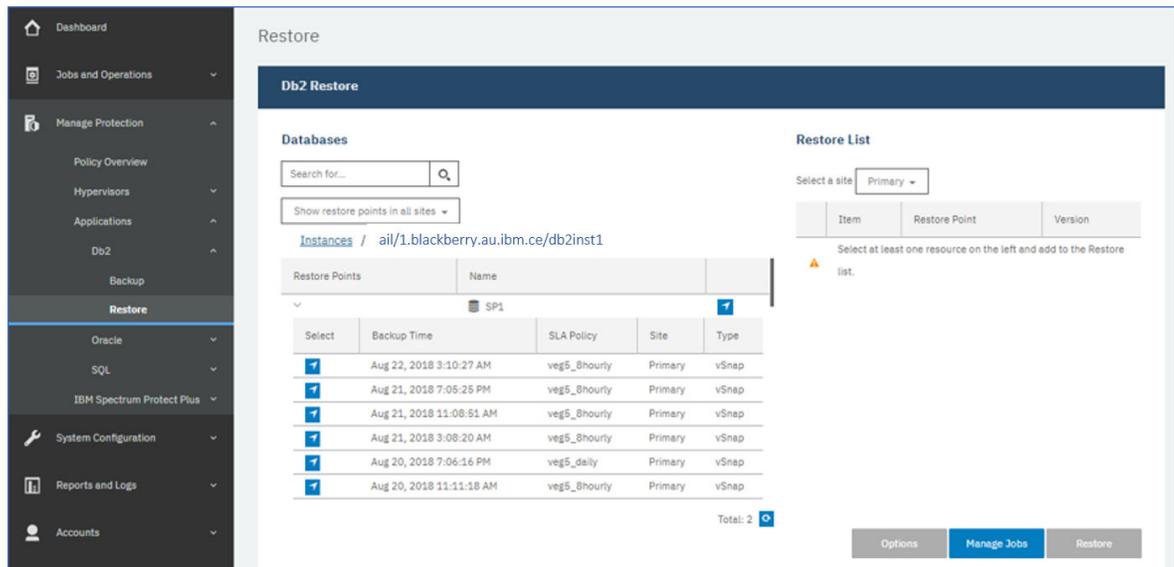


Figure 21: Selecting to add a specific backup the restore list

To remove the restore point from the list, click the delete icon .



Attention: Ensure that you review the selected options before clicking **Restore** because data will be overwritten when the **Overwrite existing data** option is selected.

5. Click **Options** to define specific options for the restore job as follows:

Restore Type is automatically set to **Production**. This is the only restore type allowed for restoring to the original instance.

6. Choose **Restore to original instance** to restore to the original server for the production restore operation.

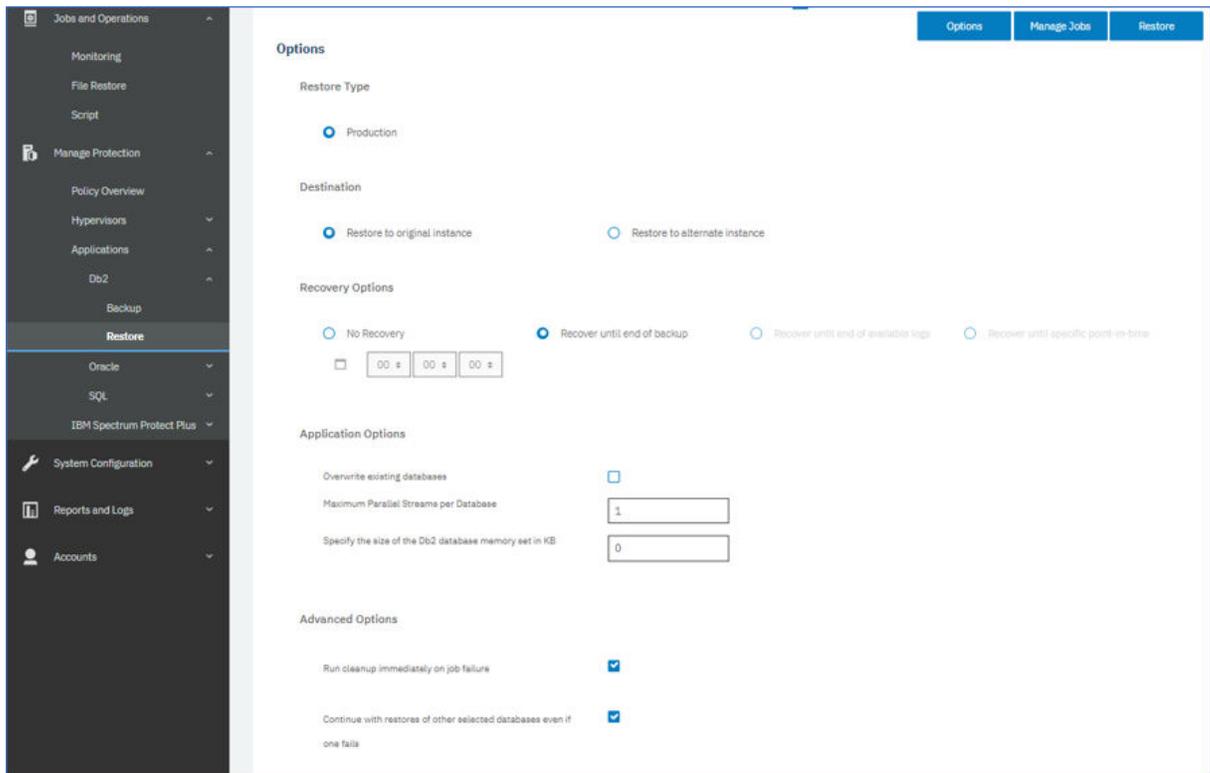


Figure 22: Restore options.

7. Edit **Recovery Options** for the restore operation.

The default value is **Recover until end of backup**.

- **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 you have enabled log backups in your Db2 backup job definition. IBM Spectrum Protect Plus uses the newest 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, it is only available when you add the overall database. When you add the overall database, the newest backup is automatically selected for the end of logs recovery.
- **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, May 8, 2018 12:18:00 AM. IBM Spectrum Protect Plus finds the restore points directly before and after the point-in-time chosen. During the recovery process, the older data backup volume and the newer log backup volume are mounted. A temporary restore point is created if the point in time is after the last backup. This recovery option is not available if you selected a specific restore point from the list. Selecting the overall database automatically selects the newest backup for the end of logs recovery process.

8. Edit **Application Options**.

- **Overwrite existing databases** this option overwrites data related to the database being restored at the target instance. If existing databases are found on the target instance, the restore operation fails.

- **Maximum Parallel Streams per Database.** If required you can choose to run the restore operation of data in parallel streams. This option is useful when 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.
9. Define advanced options for the restore job in the **Advanced Options** section as follows:
 - **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.
 10. Define **Scripts Settings** as follows:
 - Select **Pre-script** to select an uploaded script and an application or script server where the pre-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **Jobs and Monitoring > Script** pane to configure scripts and script servers.
 - Select **Post-script** to select an uploaded script and an application or script server where the post-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **Jobs and Monitoring > Script** pane to configure scripts and script servers.
 - Select **Continue job/task on script error** to continue running the job when the script associated with the job fails. When this option is enabled, in the event that a script completes with a nonzero return code, the backup or restore job continues to run and the pre-script task status returns COMPLETED. If a post-script completes with a nonzero return code, the Post-script task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task status returns FAILED.
 11. Review the job definition and click **Save**.
 12. To run the job immediately, click **Restore**. To specify a schedule for a repeated restore operation, click **Manage Jobs** to define a trigger for the job.

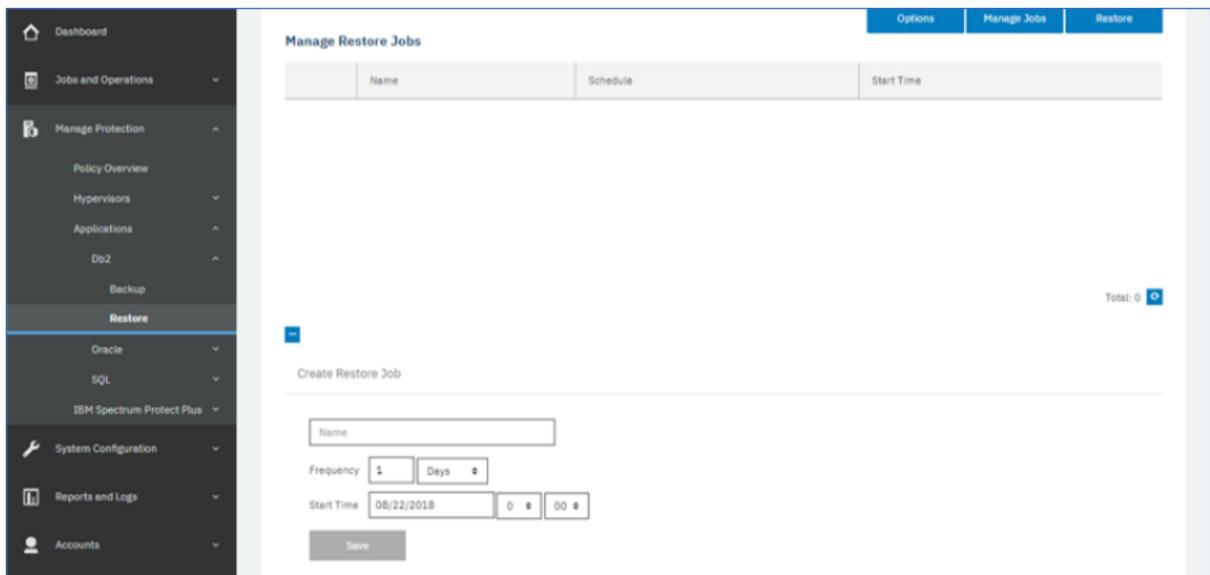


Figure 23: Restore job definition in Manage Restore Jobs

Restoring Db2 databases to an alternate instance

Select a Db2 database backup and restore that database to a Db2 instance on an alternate 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.

Before you begin

In the **Restore** pane, use the filters to show the instances from the primary or secondary sites. The default shows the restore points from all sites. If you cannot find a specific database, search by entering a database name in the search field.

Before you create a restore job for Db2, ensure that the following requirements are met.

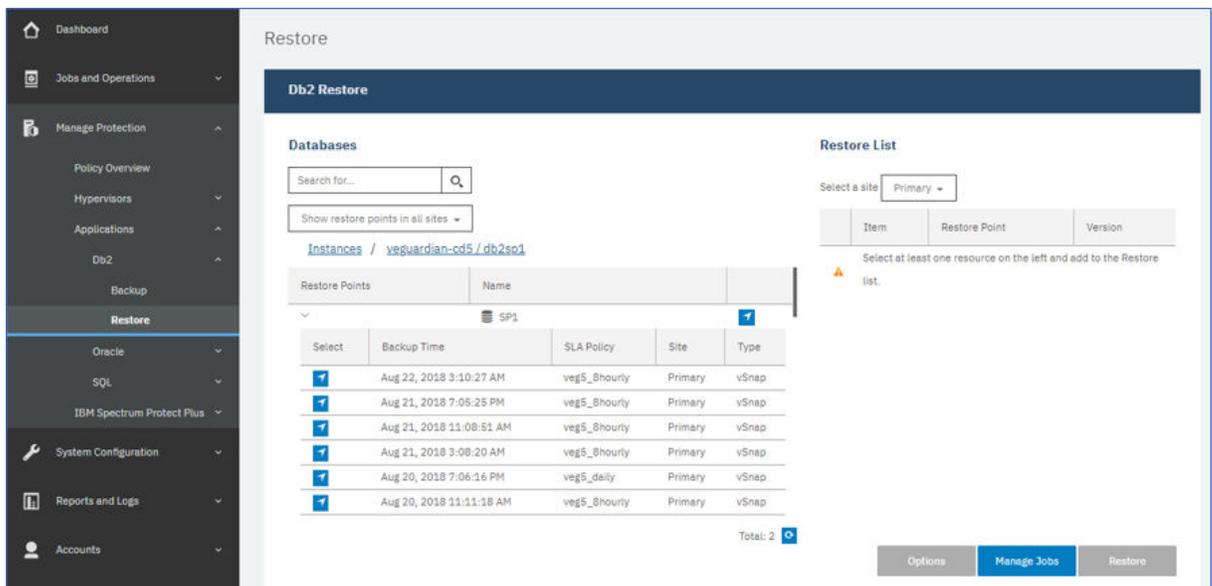
- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up Db2 data” on page 106](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see [Chapter 13, “Managing user access,” on page 169](#).

Before you start a restore operation to an alternate instance, ensure that the same 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 the Db2](#).

Tip: If data exists on the local database directory you are restoring to, and the **Overwrite existing databases** option is not selected, the restore to alternate instance fails. No other data can share the local database directory as the database to be restored. When the **Overwrite existing databases** option is selected, any existing data is wiped and the local database directory is cleaned up on the alternate host.

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > Db2 > Restore**.
2. In the **Db2 Restore** pane, click a Db2 instance to show the databases in that instance.
3. Expand the database that you want to restore to show the available restore points for that database.



The screenshot displays the 'Restore' panel for Db2. On the left is a navigation sidebar with 'Restore' selected under 'Applications > Db2'. The main area is titled 'Restore' and 'Db2 Restore'. It features a search bar, a dropdown for 'Show restore points in all sites', and a breadcrumb for 'Instances / veguardian-cd5 / db2sp1'. Below this is a table of 'Restore Points' with columns for 'Select', 'Backup Time', 'SLA Policy', 'Site', and 'Type'. The table contains six rows of data. To the right is the 'Restore List' section, which includes a 'Select a site' dropdown set to 'Primary' and a table with columns 'Item', 'Restore Point', and 'Version'. A warning icon and text prompt the user to 'Select at least one resource on the left and add to the Restore list.' At the bottom right, there are buttons for 'Options', 'Manage Jobs', and 'Restore'.

Select	Backup Time	SLA Policy	Site	Type
<input checked="" type="checkbox"/>	Aug 22, 2018 3:10:27 AM	veg5_hourly	Primary	vSnap
<input checked="" type="checkbox"/>	Aug 21, 2018 7:05:25 PM	veg5_hourly	Primary	vSnap
<input checked="" type="checkbox"/>	Aug 21, 2018 11:08:51 AM	veg5_hourly	Primary	vSnap
<input checked="" type="checkbox"/>	Aug 21, 2018 3:08:20 AM	veg5_hourly	Primary	vSnap
<input checked="" type="checkbox"/>	Aug 20, 2018 7:06:16 PM	veg5_daily	Primary	vSnap
<input checked="" type="checkbox"/>	Aug 20, 2018 11:11:18 AM	veg5_hourly	Primary	vSnap

Figure 24: Restore panel for Db2

Restore points are listed with the backup date and time, SLA policy, site information, and type.

4. Click **Options** and first choose **Destination > Restore to alternate instance**.
5. Choose a **Restore Type**.
 - – **Test**: In this mode, the agent creates a new database by using the data files directly from the vSnap repository. This option is available only when restoring to an alternate instance.
 - – **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. This restore type is the only option available when you restore data to the original instance.
 - – **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. This option is available only when restoring to an alternate instance.
6. Select the target instance that you want to restore the data to.
 - The original instance is not selectable as you cannot overwrite the original data when you select **Restore to alternate instance**.
 - Instances on different version levels cannot be selected.
 - Other instances on the same host as the original instance, cannot be selected.
 - *Original instance* refers to the instance that the backup originated from.

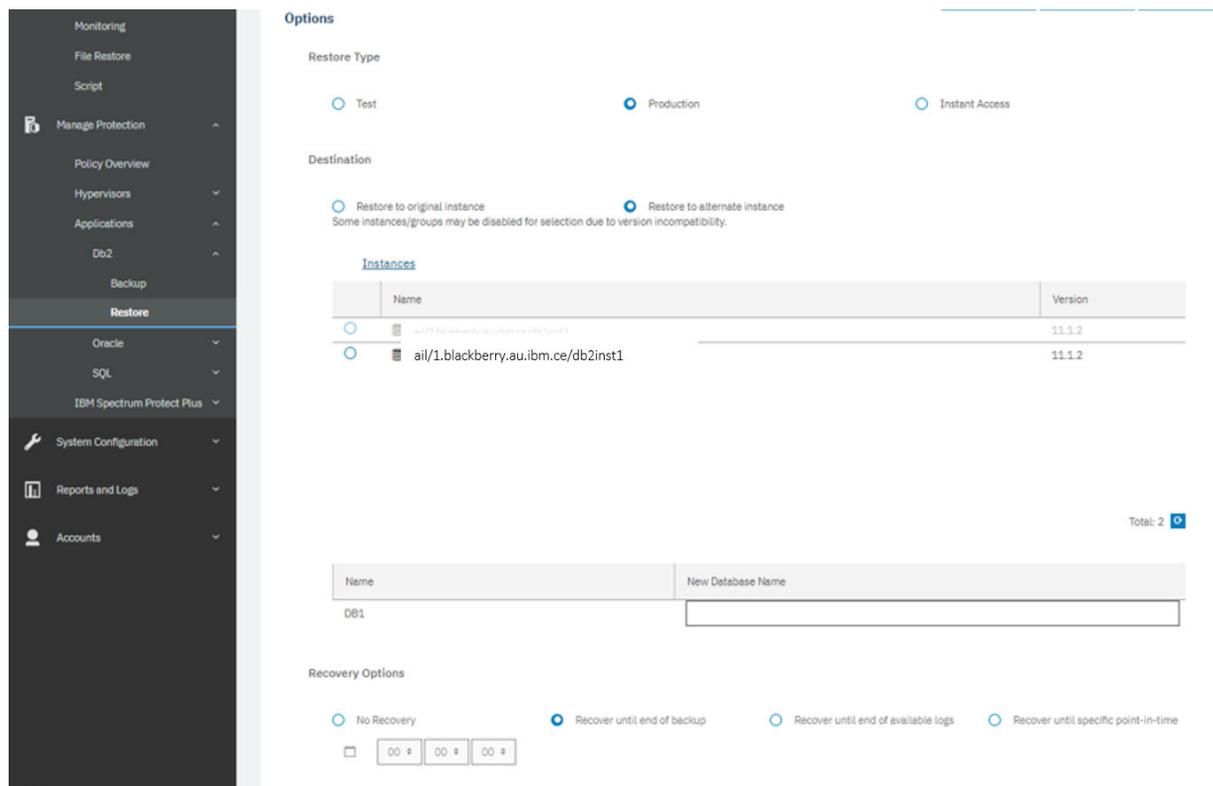


Figure 25: Restoring to alternate instance and Instances pane

7. Optional: Enter a new database name in the **New Database Name** field.
8. Edit **Recovery Options** for the restore operation.

The default value is **Recover until end of backup**.

 - **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 you have enabled log backups in your Db2 backup job definition. IBM Spectrum Protect Plus uses the newest 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, it is only available when you add the overall database. When you add the overall database, the newest backup is automatically selected for the end of logs recovery.
- **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, May 8, 2018 12:18:00 AM. IBM Spectrum Protect Plus finds the restore points directly before and after the point-in-time chosen. During the recovery process, the older data backup volume and the newer log backup volume are mounted. A temporary restore point is created if the point in time is after the last backup. This recovery option is not available if you selected a specific restore point from the list. Selecting the overall database automatically selects the newest backup for the end of logs recovery process.

9. Select application options in the **Application Options** section as follows:

- **Overwrite existing databases.** This option is selected by default during a Test restore operation to replace existing databases with the same name encountered by the restore process during recovery. If this option is not selected, the restore job fails when databases with the same name are found during the restore process. If you select this option, ensure that the Db2 log directory and the Db2 mirror log directory are clear of data.



Attention: Ensure that no other databases share the same local database directory as the original database or that data is overwritten when this choice is selected.

- **Maximum Parallel Streams per Database.** If required you can choose to run the restore operation of data in parallel streams. This option is useful when 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. Define advanced options for the restore job in the **Advanced Options** section as follows:

- **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.

11. Define **Scripts Settings** as follows:

- Select **Pre-script** to select an uploaded script and an application or script server where the pre-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **Jobs and Monitoring > Script** pane to configure scripts and script servers.
- Select **Post-script** to select an uploaded script and an application or script server where the post-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **Jobs and Monitoring > Script** pane to configure scripts and script servers.
- Select **Continue job/task on script error** to continue running the job when the script associated with the job fails. When this option is enabled, in the event that a script completes with a nonzero return code, the backup or restore job continues to run and the pre-script task status returns

COMPLETED. If a post-script completes with a nonzero return code, the Post-script task status returns COMPLETED. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task status returns FAILED.

12. Review the job definition and click **Save**.
13. To run the job immediately, click **Restore**. To specify a schedule for a repeated restore operation, click **Manage Jobs** to define a trigger for the job.

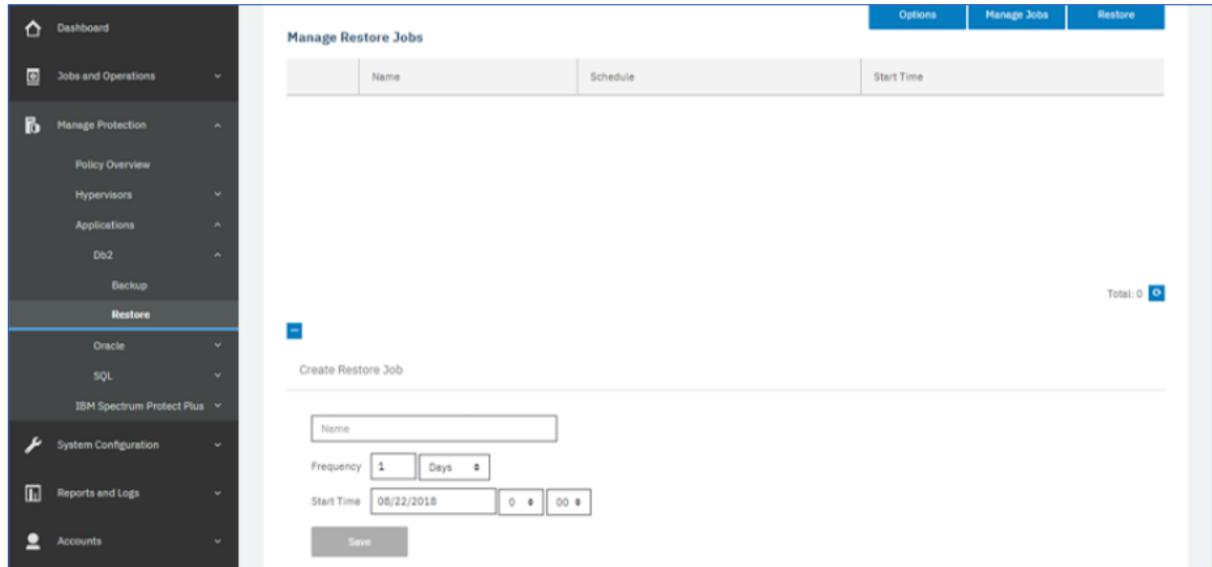


Figure 26: Restore job definition in Manage Restore Jobs

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 requirements”](#) on page 30.

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.

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.

Adding a SQL Server application server

When a 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 a 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 174 .

Related concepts

[“Managing user access” on page 169](#)

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 125](#)

Use a backup job to back up SQL Server environments with snapshots.

[“Restoring SQL Server data” on page 127](#)

Use a restore job to restore SQL Server environments from snapshots. Your SQL Server clones can be utilized and consumed instantly through IBM Spectrum Protect Plus Instant Disk Restore jobs. 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 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 new vSnap volume and creates an NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus agent mounts the share on the SQL Server server where the backup is to be completed.

When the backup is complete, 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 13, “Managing user access,” on page 169](#).
- 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](#).
- Avoid configuring log backup for a single SQL database through 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 it is backed up by the next job. This might cause point-in-time restore jobs to fail.
- IBM Spectrum Protect Plus does not support log backup of Simple recovery models.
- An Always On of a replica of a SQL Server cluster instance is not supported. Replicas are limited standalone SQL Server servers and instances.
- Failover of a 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](#).
- When a log backup of a secondary SQL Always On database fails with the following error, the backup preference of the availability group must be changed to Primary:

```
Log backup for database 'DatabaseName' on a secondary replica failed because a synchronization point could not be established on the primary database.
```

Changing the preference to Primary will back up the log from the primary replica. After a successful log backup of the primary replica is completed, the backup preference can be changed.

Take the following actions:

- Register the providers that you want to back up. For more information, see [“Adding a SQL Server application server” on page 123](#).
- Configure SLA policies. For more information, see [“Create backup policies” on page 52](#).
- Before setting up and running SQL backup jobs you must configure the Shadow Copy storage settings for the volumes where your SQL databases are located. This setting is configured once per volume. If new databases are added to the job, the setting must be configured for any new volumes containing 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 depending on the source volume size and I/O

activities, then click **OK**. The shadow copy storage area must be on the same volume or another available volume during the time of backup.

Procedure

To define a SQL backup job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Applications > SQL > Backup**.
2. Select a 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 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, or you can manually run the job by clicking **Jobs and Operations > Monitoring** and starting the job from the **Job Monitor** pane.
5. To edit options before you create the job definition, click **Select Options**. Set the job definition options.

Enable Log Backup

Select to enable IBM Spectrum Protect Plus to back up transaction logs and then protect the underlying disks.

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 by IBM Spectrum Protect Plus and must be managed separately.

When SQL backup job completes with log backups enabled, all transaction logs up to the point of the job completing are purged from the SQL Server server. Log purging occurs only if the SQL Backup job completes successfully. If log backups are disabled during a rerun of the job, log purging does not occur.

If a source database is overwritten, all old transaction logs up to that point are placed in a “condense” directory once the restoration of the original database completes. When the next run of the SQL Backup job completes, the contents of the condense folder is removed.

To complete log backups, the SQL Server agent service user must be a local Windows administrator and must have the sysadmin permission enabled to manage SQL Server agent jobs. The agent will use that administrator account to enable and access log backup jobs. The IBM Spectrum Protect Plus SQL Server agent service user must also be the same as the SQL Server service and SQL Server agent service account for every SQL Server instance to be protected.

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.

When this option is selected, point-in-time restore options are available for SQL restore operations.

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**.
The job runs as defined by your SLA policy, or can be run manually from the Job Monitor pane.
7. To configure additional options, click the **Options** field that is associated with the job in the **SLA Policy Status** section. If no additional options are configured for the job, **Not Configured** is shown in the field. Set the additional job options:

Pre-scripts and post-scripts

Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. 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 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 **Jobs and Monitoring > Script** pane.

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 Base 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 a SQL Restore job definition.	See “Restoring SQL Server data” on page 127 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 144](#)

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 **Scripts** pane, and then applied to job definitions.

Related tasks

[“Starting jobs” on page 144](#)

You can run any job on demand, even if the job is set to run on a schedule.

Restoring SQL Server data

Use a restore job to restore SQL Server environments from snapshots. Your SQL Server clones can be utilized and consumed instantly through IBM Spectrum Protect Plus Instant Disk Restore jobs. IBM Spectrum Protect Plus catalogs and tracks all cloned instances.

Before you begin

Note the following procedures and considerations before creating a restore job definition:

- Create and run a SQL backup job. For more information, see [“Backing up SQL Server data” on page 125](#).
- Review the SQL Server system requirements in [“Microsoft SQL Server requirements” on page 30](#).

- 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 13, “Managing user access,” on page 169](#).
- When completing a production restore to a SQL Server failover cluster, the root volume of the alternate 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.
- A restore to an NTFS or FAT compressed volume is not supported because of SQL Server database restrictions. For more information see [Description of support for SQL Server databases on compressed volumes](#).
- When completing 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.
- When restoring to an alternate 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 restoring to a primary instance in a SQL Always On Availability Group environment, the database is added to the target Always On database group. After the primary restore, the secondary database is seeded by the SQL server in environments where automatic seeding is supported (SQL 2016 and later). The database is then enabled on the destination availability group. The synchronization time depends on the amount of data being transferred and the connection between the primary and secondary replicas.

If automatic seeding is not supported or is disabled, a secondary restore from the restore point with the shortest LSN gap of the primary instance must be completed. Log backups with the latest point-in-time restore point created by IBM Spectrum Protect Plus must be restored if the log backup was enabled on the primary instance. The secondary database restore is completed in a restoring state and you must issue the T-SQL command to add the database to the target group. For more information, see [Transact-SQL Reference \(Database Engine\)](#).

About this task

Instant Disk Restore leverages iSCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases 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 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 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 spins up the new database using the restored files.

Procedure

To define a SQL restore job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Applications > SQL > Restore**.
2. In the **Restore** pane, review the available restore points of your SQL Servers servers.
3. 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**.
4. To select the latest restore point, complete the following actions:
 - a) Click the add to restore list icon  at the resource level.

- b) From the **Restore from site** list, select the site associated with the backup storage server that you want to recover from.

To restore a specific restore point from a specific site, complete the following actions:

- a) Expand a resource in the **Restore** pane.
- b) Click the add to restore list icon  that is associated with the restore point.

Adding the latest restore points and specific restore points to the **Restore List** is not supported. Click the remove icon  to remove restore points from the **Restore List**.

5. To run the job now using default options, click **Restore**. To schedule the job to run using default options, click **Manage Jobs** and define a trigger for the job definition.
6. To edit options before you create the job definition, click **Options**. Set the job definition options.

Destination

Set the restore destination.

Restore to original instance

Select to restore to the original instance.

Restore to primary instance

When performing restore operations in a SQL Always On environment, select to restore the database to the primary instance of the Always On Availability Group and add the database back to the group

Restore to alternate instance

Select to restore to a local destination different from the instance, then select the alternate location from available servers. When performing restore operations in a SQL Always On environment using test mode, the source availability database is restored to the selected target instance.

When performing restore operations in a SQL Always On environment using 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 destination availability group has the automatic seeding option enabled, the secondary database file paths are synchronized with the primary database. If the primary database log is not truncated, the secondary database may be added to the availability group by SQL.

New Database Name

Click the **New Database Name** field to enter an optional alternate name for the database.

Restore Type

Set the SQL Restore job to run in test, production, or Instant Access mode by default. Once the job is created, it can be run in test, production, or Instant Access mode through the **Job Sessions** pane.

Recovery Options

Set the recovery options:

No Recovery

Sets 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. The **No Recovery** option does not support production mode restores to SQL Always On groups.

Recover until end of backup

Restore the selected database to the state at the time the backup was created.

Recover until specific point in time

When log backup is enabled through a SQL backup job definition, point-in-time restore options will be available when creating a SQL 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 by a specific date and time.
- **By Transaction ID.** Select this option to configure a point-in-time recovery by transaction ID.

In a standalone 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. A temporary restore point is created if the point-in-time is after the last backup operation.

When performing restore operations in a SQL Always On environment using test mode, the restored database will join the instance where the availability group resides.

When performing restore operations in a SQL Always On environment using production mode, the restored Primary database will be joined to the availability group. If the destination availability group has the automatic seeding option enabled, the secondary database file paths are synchronized with the primary database. If the primary database log is not truncated, the secondary database may be added to the availability group by SQL.

Application Options

Set the application options:

Overwrite existing database

Enable to allow the restore job to overwrite the selected database. By default, this option is disabled.

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. Note that 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 restoring 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

Enable to automatically clean up allocated resources as part of a restore if the recovery fails.

Allow session overwrite

Select this option to replace an existing database with 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/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 encounters an existing 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 disabled, 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. Available protocols are **iSCSI** and **Fibre Channel**.

Script Settings

Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. 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 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 **Jobs and Monitoring > Script** pane.

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.

7. Click **Save**.

8. To run the job now, click **Restore**. To schedule the job click **Manage Jobs** and define a trigger for the job definition.

Related concepts

[“Configuring scripts for backup and restore operations” on page 144](#)

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 **Scripts** pane, and then applied to job definitions.

Related tasks

[“Adding a SQL Server application server” on page 123](#)

When a 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 125](#)

Use a backup job to back up SQL Server environments with snapshots.

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 requirements” on page 27](#).

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**.
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 application server.

Username

Enter your user name for the application 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.

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 application server.

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 Oracle application server, complete the following action:

Action	How to
Assign user permissions to the application server.	See “Creating a role” on page 174 .

Related concepts

[“Managing user access” on page 169](#)

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 133](#)

Use a backup job to back up Oracle environments with snapshots.

[“Restoring Oracle data” on page 135](#)

Use a restore job to restore Oracle environments from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version selected during the job definition creation and creates an NFS share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle server where the restore is to be performed. In the case of Oracle RAC, the restore is performed 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 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

During the initial base backup, IBM Spectrum Protect Plus creates a new vSnap volume and creates 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 RAC, the backup is completed from any one node in the cluster. When the backup is complete, the IBM Spectrum Protect Plus agent unmounts the share from the Oracle 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 13, “Managing user access,”](#) on page 169.
- 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 datafiles 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:

- Register the providers that you want to back up. For more information, see [“Adding an Oracle application server”](#) on page 131.
- Configure SLA policies. For more information, see [“Create backup policies”](#) on page 52.

Procedure

To define an Oracle backup job, complete the following steps:

1. In the navigation pane, click **Manage Protection > Applications > Oracle > Backup**.
2. Select Oracle homes, databases, and ASM diskgroups to back up. Use the search function to search for available instances.
3. Click **Select 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, or you can manually run the job by clicking **Jobs and Operations > Monitoring** and starting the job from the **Job Monitor** pane.
5. To edit options before you create the job definition, click **Select Options**. Set the job definition options.

Enable Log Backup

Select to enable IBM Spectrum Protect Plus to automatically create a log backup volume and mount it to the application server. IBM Spectrum Protect Plus then uses cron to configure a scheduled job that completes a transaction log backup to that volume at the frequency specified through the **Frequency** setting.

For Oracle Real Application Clusters (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. After a successful database backup, older logs are deleted automatically from this log backup volume.

IBM Spectrum Protect Plus does not manage the retention of other archived log locations. Administrators must continue to manage those logs using their existing log retention policies.

Select **Truncate source logs after successful backup** to automatically delete older archived logs that are no longer needed from the log backup destination and primary destination after every successful database backup. If the option is cleared, archived logs on the primary log destination are not deleted. Set the log retention through the **Primary log retention in days** setting.

When this option is selected, point-in-time restore options are available for Oracle restore operations.

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 **Options** field that is associated with the job in the **SLA Policy Status** section. If no additional options are currently configured for the job, **Not Configured** is shown in the field. Set the additional job options, and then click **Save**:

Pre-scripts and Post-scripts

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 **Jobs and Monitoring > Script** pane.

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 Base 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 135 .

Related concepts

[“Configuring scripts for backup and restore operations” on page 144](#)

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 **Scripts** pane, and then applied to job definitions.

Restoring Oracle data

Use a restore job to restore Oracle environments from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version selected during the job definition creation and creates an NFS share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle server where the restore is to be performed. In the case of Oracle RAC, the restore is performed on all nodes in the cluster.

Before you begin

Note the following procedures and considerations before creating a restore job definition:

- Create and run an Oracle backup job. For more information, see [“Backing up Oracle data” on page 133](#).
- Review the Oracle system requirements in [“Oracle requirements” on page 27](#)
- 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 13, “Managing user access,” on page 169](#).
- 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.
- If an Oracle database is mounted but not opened during a backup job, IBM Spectrum Protect Plus cannot determine the database tempfile settings related to autoextensibility and maximum size. When a database is restored from this restore point, IBM Spectrum Protect Plus cannot recreate 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.

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 using the files in the vSnap volume.

Test mode

In test mode, the agent creates a new database 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 spins up the new database 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 > Restore**.
2. In the **Restore** pane, review the available restore points of your Oracle instances.
3. Use the search function to search for available instances and toggle the displayed instances through the **View** filter.
4. To select the latest restore point, complete the following actions:
 - a) Click the add to restore list icon  at the resource level.
 - b) From the **Restore from site** list, select the site associated with the backup storage server that you want to recover from.

To restore a specific restore point from a specific site, complete the following actions:

- a) Expand a resource in the **Restore** pane.
 - b) Click the add to restore list icon  that is associated with the restore point.
- Adding the latest restore points and specific restore points to the **Restore List** is not supported. Click the remove icon  to remove restore points from the **Restore List**.
5. To run the job now using default options, click **Restore**. To schedule the job to run using default options, click **Manage Jobs** and define a trigger for the job definition.
 6. To edit options before you create the job definition, click **Options**. Set the job definition options.

Destination

Set the restore destination.

Restore to original location

Select to restore to the original server.

Restore to alternate location

Select to restore to a local destination different from the server, then select the alternate location from available servers.

New Database Name

Click the **New Database Name** field to enter an optional alternate name for the database.

Restore Type

Set the Oracle restore job to run in test, production, or Instant Access mode by default. Once the job is created, it can be run in test, production, or Instant Access mode through the **Job Sessions** pane.

Recovery Options

Set the recovery options:

Recover until end of backup

Restore the selected database to the state at the time the backup was created.

Recover until specific point in time

When log backup is enabled through an Oracle Backup job definition, point-in-time restore options will be available when creating 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 by 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. A temporary restore point is created if the point-in-time is after the last backup.

Application Options

Set the application options:

Overwrite existing database

Enable to allow the restore job to overwrite the selected database. By default, this option is disabled.

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. Note that 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 restoring an Oracle database to its original location using its original database name.

Init Params

This option controls the initialization parameters that are used to start up the recovered database in Oracle Test and Production workflows

Source. This is the default option. 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.
- The `db_name` and `db_unique_name` are updated to reflect the new name of the database if a new name is specified.
- 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 containing the initialization parameters that IBM Spectrum Protect Plus should use.

The specified path must be 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 form `name = value`. Comments beginning with the `#` character are ignored.

IBM Spectrum Protect Plus reads the template pfile and copies the entries to the new pfile that will be used to start up 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 to automatically clean up allocated resources as part of a restore if the recovery fails.

Allow session overwrite

Select this option to replace an existing database with 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/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 encounters an existing 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 disabled, 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. Available protocols are **iSCSI** and **Fibre Channel**.

Script Settings

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 **Jobs and Monitoring > Script** pane.

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.

7. Click **Save**.

8. To run the job now, click **Restore**. To schedule the job click **Manage Jobs** and define a trigger for the job definition.

Related concepts

[“Configuring scripts for backup and restore operations” on page 144](#)

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 **Scripts** pane, and then applied to job definitions.

Related tasks

[“Adding an Oracle application server” on page 131](#)

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.

Chapter 9. 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, search data, and job information are backed up to a vSnap server 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, search data, and imported keys and certificates to a vSnap server that is defined in the associated SLA policy.

Before you begin

Ensure that the appropriate SLA policy is available. To optimize backup jobs, create SLA policies specifically for backing up IBM Spectrum Protect Plus. To create an SLA policy, see [“Creating an SLA policy”](#) on page 67.

Restriction: You cannot select the onboard vSnap server as the target of the IBM Spectrum Protect Plus backup SLA policy. 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 creating an SLA policy for backup

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. The SLA policy defines the scheduling of the catalog backup, along with the backup destination, replication, and offloading settings.
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 > Monitoring** and starting the job from the **Job Monitor** pane. For instructions, see [“Start a job”](#) on page 59.

Restoring the IBM Spectrum Protect Plus application

Restore IBM Spectrum Protect Plus configuration settings, restore points, search data, 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.

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.

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.

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 offload relationship. Run the replication or offload-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 restoration point to expire:

1. In the navigation pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore Point Retention**.
2. To search for a restore point by name, enter a search string. For more information about using the search function, see [Appendix A, "Search guidelines," on page 183](#).
3. 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 **Actions > Expire** for the restore point that you want to expire.
6. To confirm the expiration, in the dialog box, click **Yes**.

Results

The restore point is removed during the next run of the maintenance job.

Related concepts

["Job types" on page 143](#)

Jobs are used to run backup, restore, maintenance, and inventory operations in IBM Spectrum Protect Plus.

Chapter 10. Managing data protection jobs

You can run jobs on demand, pause or cancel running jobs, and pause scheduled jobs.

To view and manage jobs, click **Jobs and Operations > Monitoring** in the navigation pane.

Jobs are displayed in the **Job Monitor** pane. The name, type, status, last run time, and next scheduled run time (if applicable) are shown for each job.

To view logs for the job, expand the job.

To download a job log as a .csv file, click the download icon  .

Job types

Jobs are used to run backup, restore, maintenance, and inventory 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 and inventory jobs are predefined and not modifiable.

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.

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. You can create a schedule to run the job or you can run the job on demand.

The job name depends on whether you run the job on demand or on a schedule. If you run a restore operation on demand, the job name onDemandRestore is auto generated.

If create a job to run on a schedule, you must specify a job name.

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.

Starting jobs

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 > Monitoring**.
2. To start the job session, click the **Actions** menu that is associated with the job, and then click **Start**.
3. Expand the running job session to view the job session details.

The following items are included in the details:

- Duration of the job
- Start time of the job
- End time of the job
- Total number of protected VMs
- Total number of failed VMs

Pausing and resuming jobs

You can pause and resume a scheduled job or a job that is running. 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 > Monitoring**.
2. To pause a job schedule, click the **Actions** menu that is associated with the job, and then click **Pause Schedule**.
3. To resume the job schedule, click the **Actions** menu that is associated with the job that you want to resume and click **Release 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 > Monitoring**.
2. To cancel a running job session, click the **Actions** menu that is associated with the job, and then click **Cancel**.

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 **Scripts** pane, 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, expand **Jobs and Operations**, and then click **Script**.
2. In the scripts section, click **Upload Script**.
The **Upload Script** pane displays.
3. Click **Browse** to select a local script to upload.
4. Click **Save**.

The script displays 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 145 .

Adding a script to a server

Add the script to a server from which it will run.

Procedure

Complete the following steps to designate a script to a server:

1. In the navigation pane, expand **Jobs and Operations**, and then click **Script**.
2. In the script server 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 11. Configuring and maintaining the IBM Spectrum Protect Plus system environment

System management tasks include adding backup storage, managing sites, and registering Lightweight Directory Access Protocol (LDAP) or Simple Mail Transfer Protocol (SMTP) servers.

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. Z File System (ZFS) update packages are also provided for vSnap stand-alone instances.



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

Managing sites

A site is a user-defined grouping of backup storage servers that is generally based on a physical or logical location to help quickly identify and interact with backup data.

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, and then click **Save**.

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

Editing a site

Revise site names 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, and then click **Save**.

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 177](#)

Add a user account for an LDAP group to IBM Spectrum Protect Plus.

[“Scheduling a report” on page 166](#)

You can schedule customized reports in IBM Spectrum Protect Plus to run at specific times.

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.

After the SMTP server is added, the **Add SMTP 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

[“Scheduling a report” on page 166](#)

You can schedule customized reports in IBM Spectrum Protect Plus to run at specific times.

Adding an LDAP server

You must add an LDAP server to create IBM Spectrum Protect Plus user accounts by using an LDAP group. These accounts allows users to access IBM Spectrum Protect Plus by using LDAP user names and passwords. Only one LDAP server can be associated with an instance of IBM Spectrum Protect Plus virtual appliance.

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 177](#)

Add a user account for an LDAP group to IBM Spectrum Protect Plus.

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 <code>serveradmin</code> password. The default password is <code>sppDP758</code> . You are prompted to change this password during the first logon.

What to do next

Review the configuration of the IBM Spectrum Protect Plus virtual appliance.

Related concepts

[“System requirements” on page 13](#)

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 173](#)

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 <code>serveradmin</code> password. The default password is <code>sppDP758</code> . You are prompted to change this password during the first login.

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. To view the current time zone, select **Product Information** from the main page of the Administrative Console.

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 152](#)

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 <code>serveradmin</code> password. The default password is <code>sppDP758</code> . You are prompted to change this password during the first logon.

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 file referenced below.

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.

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.

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.

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 interface

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 prompt. Issue the following command:

```
# sudo bash
```

2. Open port 9000 on the firewall by issuing the following command:

```
# firewall-cmd --add-port=9000/tcp
```

3. Run the tool by issuing the following command:

```
# java -Dserver.port=9000 -jar /opt/ECX/spp/public/assets/tool/ngxddd.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, 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.

6. Click **Save**.
7. 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.
8. Stop the tool by issuing the following command on the command line:

```
ctl-c
```

9. Protect your storage environment by resetting the firewall. Issue the following commands:

```
# firewall-cmd --zone=public --remove-port=9000/tcp
# firewall-cmd --runtime-to-permanent
# firewall-cmd --reload
```

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. Copy the .jar file to a target remote server with Java installed, and then launch the tool from a command-line interface.

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 `pvdisplay` and `vgdisplay` commands.

Adding a disk to the virtual appliance

Use the vCenter client to edit the settings of the virtual machine.

Before you begin

To run commands, you must connect to the command line for the IBM Spectrum Protect Plus virtual appliance by using Secure Shell (SSH) and log in with the user ID `serveradmin`. The default initial password is `sppDP758` and you are prompted to change the password when you log on for the first time.

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 command:

```
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` and you are prompted to change the password when you log in for the first time.

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 `pvdisplay` command.
4. To create a new Physical Volume (PV), issue the `pvcreate /dev/sdd1` command.
5. To view the new PV from `/dev/sdd1`, issue the `pvdisplay` command.
6. To review the Volume Group (VG), issue the `vgdisplay` command.
7. To add the Physical Volume (PV) to the Volume Group (VG) and increase the space of the VG, issue the following command:

```
vgextend data_vg /dev/sdd1
```

8. To verify that `data_vg` is extended, and that free space is available for logical volumes (or `/data` volume) to use, issue the `vgdisplay` command.
9. To review the Logical Volume (LV) `/data` volume, issue the `lvdisplay` command. The usage of the `/data` volume displays.
10. To add the space of the LV `/data` volume to the total volume capacity, issue the `lvextend` command.
In this example, 20 GB of space is being added to a 100 GB volume.

```
[serveradmin@localhost ~]# lvextend -L120gb -r /dev/data_vg/data
Size of logical volume data_vg/data changed from 100.00 GiB to 120.00 GiB .
Logical volume data successfully resized
resize2fs 1.41.12 (date)
Filesystem at /dev/mapper/data_vg-data is mounted on /data; on-line
resizing required
old desc_blocks = 7, new_desc_blocks = 8
Performing an on-line resize of /dev/mapper/data_vg-data to 31195136
(4k) blocks.
The filesystem on /dev/mapper/data_vg-data is now 31195136 blocks
long.
```

After you run the preceding command, the size of the `/data` volume is displayed in `lvdisplay` command output as 120 GB:

```
[serveradmin@localhost ~]# lvdisplay
--- Logical volume ---
LV Path: /dev/data_vg/data
LV Name: data
VG Name: data_vg
LV UUID: [uuid]
LV Write Access: read/write
LV Creation host, time localhost.localdomain, [date, time]
LV Status: available
# open: 1
LV Size: 120.00 GiB
Current LE: 30208
Segments : 2
Allocation inherit
Read ahead sectors: auto
- currently set to: 256
Block device: 253:1
[serveradmin@localhost ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/sda3 14G 2.6G 11G 20% /
tmpfs 16G 0 16G 0% /dev/shm
/dev/sda1 240M 40M 188M 18% /boot
/dev/mapper/data_vg-data
118G 6.4G 104G 6% /data
/dev/mapper/data2_vg-data2
246G 428M 234G 1% /data2
```

Chapter 12. 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:

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 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

[“Types of reports” on page 161](#)

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 VMs report

Run the Protected VMs report to view the protection status of your virtual machines. You can modify the report to show unprotected virtual machines and display 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 include unprotected virtual machines in the report, select **Show Unprotected VMs**. 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 Databases report

Run the Protected Databases report to view the protection status of your databases. You can modify this report to show unprotected databases, which will display the total number of databases added to the IBM Spectrum Protect Plus inventory before running backup jobs.

Use the report options to filter by database type and specific databases to display. To include unprotected databases in the report, select **Show Unprotected Databases**. To view databases that are protected through hypervisor-based backup jobs, select **Show Databases Protected as part of Hypervisor Backup**.

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.

Related concepts

[“Types of reports” on page 161](#)

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.

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 161](#)

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

VM environment reports

IBM Spectrum Protect Plus provides virtual machine environment reports to display the storage utilization and status of your virtual machines and datastores.

To view virtual machine environment reports, complete the following steps:

1. In the navigation pane, click **Reports and Logs > Reports**.
2. Expand **VM Environment** in the **Reports** pane.

The following reports are available:

VM Datastores report

Review the storage utilization of your datastores, including the total free space, provisioned space, and capacities. Run this report to view your datastores, the number of virtual machines on the datastores, and the percentage of space available. Use the report options to filter by hypervisor type and to select specific hypervisors to display. The **Detail View Filter** controls the datastores to display in the **Detail View** based on the percentage of space used. Use the **Show Only Orphaned Datastores** filter to view datastores that do not have any virtual machines assigned to them, or virtual machines that are in an inaccessible state. The reason for a datastore to be in an orphaned state is displayed in the **Datastore** field in the **Detail View**.

VM LUNs report

Review the storage utilization of your virtual machine logical unit numbers (LUNs). Run this report to view your LUNs, associated datastores, capacities, and storage vendors. Use the report options to filter by hypervisor type and to select specific hypervisors to display. Use the **Show Only Orphaned Datastores** filter to view datastores that do not have any virtual machines assigned to them, or virtual machines that are in an inaccessible state.

VM Snapshot Sprawl report

This report displays the age, name, and number of snapshots that are used to protect your Hypervisor resources. Use the report options to filter by hypervisor type and to select specific hypervisors to display. Use the **Snapshot Creation Time** filter to display snapshots from specific periods of time.

VM Sprawl report

Review the status of your virtual machines, including virtual machines that are powered off, powered on, or suspended. Run this report to view unused virtual machines, the date and time when they were powered off, and virtual machine templates. Use the report options to filter by hypervisor type and to select specific hypervisors to display. The report can be further filtered by power state over time,

including Days Since Last Powered Off and Days Since Last Suspended. The **Quick View** section displays a pie chart of used and free space on your virtual machines based on power state. Use the **Hypervisor** filter to display virtual machines on all hosts or a specific host. Information in the **Detail View** is categorized by power state. A separate table is provided for virtual machine templates.

VM Storage report

Review your virtual machines and associated datastores in this report. View associated datastores and provisioned space of the datastores. Use the report options to filter by hypervisor type and to select specific hypervisors to display. The **Detail View** displays associated datastores and the amount of space on the datastore that is allocated for virtual disk files.

Related concepts

[“Types of reports” on page 161](#)

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.

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, click **Options**, set the report parameters, 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 161](#)

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 a 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 161](#)

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 customized 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. Select a report type.
3. Select the report that you want to schedule.
4. Click **Options** to edit the report parameters.
5. Enter values in the **Name** and **Description** fields for the report.
6. Set the parameters for the report.
7. In the **Schedule Report** section, click **Define Schedule**.
8. Define a trigger for the report.
9. Enter an address to receive the scheduled report in the email field, and then click **Add a recipient**.
10. 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 161](#)

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 177](#)

Before a user can log on to IBM Spectrum Protect Plus and use the available functions, a user account must be created in IBM Spectrum Protect Plus.

Chapter 13. 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 hypervisors 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.

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 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.

The user account `admin` is used to set up IBM Spectrum Protect Plus. You cannot modify the credentials for this user other than to change to the password. You cannot delete the account. This account is assigned to the `SUPERUSER` role, which has access to all functions of the product.

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.

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 171.</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 174.

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.
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.

Resource Type	Subtype	Description
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 affect 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 a user is logged in when their permissions or access rights are changed, the user must log out and log in again for the updated permissions to take affect.
- You can edit any resource group that is not designated as **Cannot be modified**.

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:

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 150](#).

Application Admin

The Application Admin role allows a user to register and modify application database resources that are delegated by an administrator, as well as associate application databases to assigned SLA policies, perform backup and restore operations, and run and schedule reports delegated by an administrator.

Access to specific application servers must be granted by an administrator through the **Accounts > Resource Groups** pane. This step is also required if the **All Resources** resource group is selected because access to application servers are not included by default.

Backup Only

The Backup Only role allows users to complete the following actions:

- Run, edit, and monitor 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**.

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.

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.

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	<p>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.</p> <p>b. Select additional permissions to apply to the role, and delete permissions that are not required.</p> <p>To view available permissions and their usage, see “Permission types” on page 175.</p>

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.
Hypervisor	Register, view, edit, deregister, options	Used to interact with hypervisor virtual machines, such as VMware or Hyper-V virtual machines.
Identity	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.
VADP Proxy	Register, view, edit, deregister	Used to interact with VADP
Report	Create, view, edit, delete	Used to interact with reports.
Resource Group	Create, view, edit, delete	Used to interact with resource groups, which define the IBM Spectrum Protect Plus resources that are made available to a user.
Role	Create, view, edit, delete	Used to interact with roles, which define the actions that can be performed on the resources defined in a resource group.
Script	Upload, view, replace, delete	Used to interact with prescripts and postscripts that are added to IBM Spectrum Protect Plus and run before or after a job.

Name	Permissions	Description
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.
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, which associated a resource group with a role, and provides 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 affect 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 a user is logged in when their permissions or access rights are changed, the user must log out and log in again for the updated permissions to take affect.
- 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 user 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 170 and [“Creating a role”](#) on page 174.

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

Add a user account for an LDAP group to IBM Spectrum Protect Plus.

Before you begin

Review the following procedures before you create a user account for an LDAP group:

- Register an LDAP provider in IBM Spectrum Protect Plus. See [“Adding an LDAP server”](#) on page 149.
- If you want to use custom roles and resource groups, create them before you create a user. See [“Creating a resource group”](#) on page 170 and [“Creating a role”](#) on page 174.

Procedure

Complete the following steps to create a user account for an LDAP group:

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. Select an LDAP group.
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.

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 a user is logged in when their permissions or access rights are changed, the user must log out and log in again for the updated permissions to take affect.

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 **Identify 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 a 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 utilizing 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.

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 14. 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**.
2. Open or save the file log zip file, which contains individual log files for different IBM Spectrum Protect Plus components.

What to do next

To troubleshoot issues, complete the following steps:

1. Analyze the log files and take appropriate actions to resolve the issue.
2. If you cannot resolve the issue, submit the log files to IBM Software Support for assistance.

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 IBM Spectrum Protect Plus

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the [Accessibility section of the IBM Knowledge Center help \(www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility\)](http://www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

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.

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, or to request an Accessibility Conformance Report (VPAT) for this product, see [IBM Accessibility \(www.ibm.com/able\)](http://www.ibm.com/able). A VPAT outlines a product's level of conformance with the revised U.S. Section 508, the EN 501 349 standard in Europe, and WCAG 2.0.

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

accessibility features [185](#)
adding
 Hyper-V servers [90](#)
 identities [179](#)
 LDAP server [149](#)
 Oracle application servers [131](#)
 sites [88](#), [147](#)
 SMTP server [148](#)
 SQL Server application servers [123](#)
 vCenter Server instances [71](#)
 virtual disks to a vCenter virtual machine [155](#)
 vSnap servers [42](#)
Adding Db2 [103](#)
Administrative Console, logging on to [150](#)
application server
 Db2 [101](#)

B

backup jobs
 creating
 Hyper-V [92](#)
 IBM Spectrum Protect Plus [141](#)
 Oracle [133](#)
 SQL Server [125](#)
 VMware [75](#)
 excluding VMDKs from [79](#)
 starting
 on demand [144](#)
 on schedule [67](#)
backup policies, *See* SLA policies

C

creating
 reports [165](#)
 resource groups [170](#)
 roles [174](#)
 SLA policies [67](#)
 users
 individual [177](#)
 LDAP group [177](#)
 VADP proxies [80](#)

D

Data Protection for VMware
 deployment checklist [8](#)
Db2
 system requirements [24](#)
deleting
 identities [179](#)
 resource groups [172](#)
 roles [176](#)

deleting (*continued*)
 sites [148](#)
 SLA policies [69](#)
 users [178](#)
disability [185](#)

E

editing
 identities [179](#)
 resource groups [172](#)
 roles [176](#)
 sites [147](#)
 SLA policies [69](#)
 users [178](#)

F

fenced network, creating [86](#)
files
 restoring [98](#)
 searching for [183](#)

H

Hyper-V
 adding [90](#)
 backup job, creating [92](#)
 installing on virtual appliance [36](#)
 restore job, creating [95](#)
 servers
 detecting resources for [91](#)
 enabling WinRM [91](#)
 testing connection to [92](#)
 virtual appliance
 accessing [153](#)

I

identities
 adding [179](#)
 deleting [179](#)
 editing [179](#)
installing
 download packages, obtaining [33](#)
 virtual appliance
 on Hyper-V [36](#)
 on VMware [34](#)
 vSnap servers
 Hyper-V environment [41](#)
 physical environment [39](#)
 VMware environment [40](#)

J

jobs

jobs (*continued*)
canceling [144](#)
names of [143](#)
pausing [144](#)
releasing [144](#)
starting
on demand [144](#)
on schedule [67](#)
types of [143](#)

L

LDAP
group, creating a user account for [177](#)
server
adding [149](#)
Linux-based vCenter virtual appliance, backing up [79](#)
logs
audit
downloading [166](#)
viewing [166](#)
system
downloading [181](#)
viewing [181](#)

N

network
testing [154](#), [155](#)
New in IBM Spectrum Protect Plus Version 10.1.2 [vii](#)

O

Oracle
application servers
adding [131](#)
detecting resources for [132](#)
testing connection to [133](#)
backup job, creating [133](#)
restore job, creating [135](#)
system requirements [27](#)

P

prerequisites
Db2 [101](#)

Q

quick start [51](#)

R

reports
custom, creating [165](#)
running
on demand [165](#)
on schedule [166](#)
types of
backup storage utilization [161](#)
protection [162](#)
system [163](#)

reports (*continued*)
types of (*continued*)
VM environment [164](#)

resource groups
creating [170](#)
deleting [172](#)
editing [172](#)
types of [171](#)
restore jobs
creating
Hyper-V [95](#)
IBM Spectrum Protect Plus [141](#)
Oracle [135](#)
SQL Server [127](#)
VMware [82](#)
running
Hyper-V [95](#)
Oracle [135](#)
SQL Server [127](#)
VMware [82](#)
restore points, managing [142](#)
roles
creating [174](#)
deleting [176](#)
editing [176](#)
permission types [175](#)

S

Schedule jobs
Backup [108](#)
scripts for backup and restore operations
uploading [145](#)
service level agreement, *See* SLA policies
sites
adding [88](#), [147](#)
deleting [148](#)
editing [147](#)
SLA [108](#)
SLA policies
adding [67](#)
deleting [69](#)
editing [69](#)
SMTP
server
adding [148](#)
SQL Server
application servers
adding [123](#)
detecting resources for [124](#)
testing connection to [124](#)
backup job, creating [125](#)
requirements for data protection [122](#)
restore job, creating [127](#)
system requirements [30](#)
SSL certificate, uploading
from administrative console [152](#)
from command line [152](#)
starting
IBM Spectrum Protect Plus [52](#)
jobs
on demand [144](#)
on schedule [67](#)
system requirements

system requirements (*continued*)
 components [13](#)
 Db2 [24](#)
 file index and restore [21](#)
 hypervisors [21](#)
 Oracle [27](#)
 SQL Server [30](#)

T

time zone, setting [151](#)

U

user access [5](#), [169](#)

users

- deleting [178](#)
- editing [178](#)
- individual, creating [177](#)
- LDAP group, creating [177](#)
- resource groups
 - creating [170](#)
 - deleting [172](#)
 - editing [172](#)
 - types of [171](#)
- roles
 - creating [174](#)
 - deleting [176](#)
 - editing [176](#)
 - permission types [175](#)

V

VADP proxies

- creating [80](#)
- options, setting [81](#)

virtual appliance

- accessing
 - in Hyper-V [153](#)
 - in VMware [153](#)
- adding a disk to [156](#)
- adding storage capacity [156](#)
- installing
 - on Hyper-V [36](#)
 - on VMware [34](#)

VMware

- backup job, creating [75](#)
- backup job, excluding VMDKs from SLA policy [79](#)
- installing on virtual appliance [34](#)
- restore job
 - creating a fenced network [86](#)
- restore job, creating [82](#)
- vCenter Server instances
 - adding [71](#)
- vCenter Server, detecting resources [75](#)
- vCenter Server, testing connection to [75](#)
- virtual appliance
 - accessing [153](#)
- virtual machine privileges, required [73](#)

vSnap server

- administering
 - network administration [49](#)
 - storage administration [47](#)

vSnap server (*continued*)

- deleting [43](#)
- editing [43](#)
- initializing
 - advanced [44](#)
 - simple [44](#)
- replication partnership, establishing [46](#)
- storage options, managing [45](#)
- storage pools, expanding [45](#)

vSnap servers

- adding [42](#)
- installing
 - Hyper-V environment [41](#)
 - physical environment [39](#)
 - VMware environment [40](#)
- uninstalling [50](#)

W

WinRM, enabling for connection to Hyper-V servers [91](#)



Printed in USA